

Top 10 Surface Pro Fixes

Windows IT Pro

A PENTON PUBLICATION

NOVEMBER 2013 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

Manage Your Servers Remotely in Windows Server 2012

Active Directory Claims for
Windows Server 2012 File
Service Access Control

PowerShell Console
Configuration

System Center 2012
Configuration
Manager Updates

Failover Clustering in
Windows Server 2012 R2

Software Developer's

new ideas & solutions for professional programmers **JOURNAL**

Check us out! We're constantly changing for you!



BUY NOW 

With our End of summer sale you can **get access to our huge archive** (nearly 60 issues) **and to all upcoming issues**. We release around **48 SDJournal magazines per year**. You can subscribe us with **30% discount**. All you need to do is type coupon code: **JJ30** while subscribing.

<http://sdjournal.org/subscription/>



Windows IT Pro Store

eLearning Classes

eBooks

On-Demand Training

In-Person Training

Posters

Videos

Plus you can **RENEW** your subscription or
UPGRADE to VIP membership while
you're there!

Stop by the store today!

WindowsITPro

COVER STORY ▼

Windows Server 2012 33

Remote Server Management

— John Savill

Windows Server 2012's new Server Manager helps organizations move to a true remote management methodology, letting you manage multiple machines simultaneously from a graphical interface running on a client desktop.

Features

45 Use Active Directory Claims for Windows Server 2012 File Service Access Control

Sean Deuby

57 PowerShell Basics: Console Configuration

Robert Sheldon

72 Using Microsoft System Center 2012 Configuration Manager for Updates

Kent Agerlund

Products

93 New & Improved

Interact

88 Ask the Experts

In Every Issue

97 Advertiser Directory

97 Directory of Services

97 Vendor Directory

Chat with Us



Facebook



Twitter



LinkedIn

Columns



6

[Need to Know](#)

Surface 2 and Surface Pro 2

Paul Thurrott



14

[Windows Power Tools](#)

PowerShell Cmdlets for DNS

Mark Minasi



17

[Top 10](#)

Top 10 Ways Microsoft Can Fix the Surface Pro

Michael Otey



20

[Enterprise Identity](#)

The New Microsoft and the Future of the Identity Professional

Sean Deuby



24

[What Would Microsoft Support Do?](#)

Failover Clustering in Windows Server 2012 R2

John Marlin

Editorial

Editorial Director: Megan Keller
Editor-in-Chief: Amy Eisenberg
Senior Technical Director: Michael Otey
Technical Director: Sean Deuby
Senior Technical Analyst: Paul Thurrott
IT Community Manager: Rod Trent
Systems Management, Networking,
Hardware: Jason Bovberg
Scripting: Blair Greenwood
SharePoint, Active Directory, Security,
Virtualization: Caroline Marwitz
SQL Server, Developer Content:
Megan Keller
Managing Editor: Lavon Peters
Editorial SEO Specialist: Jayleen Heft

Senior Contributing Editors

David Chernicoff, Mark Minasi,
Tony Redmond, Paul Robichaux,
Mark Russinovich, John Savill

Contributing Editors

Alex K. Angelopoulos, Michael Dragone,
Jeff Felling, Brett Hill, Dan Holme,
Darren Mar-Elia, Eric B. Rux,
William Sheldon, Curt Spanburgh,
Bill Stewart, Orin Thomas, Douglas Toombs,
Ethan Wilansky

Art & Production

Senior Graphic Designer: Matt Wiebe
Director of Production: Dylan Goodwin
Group Production Manager:
Julie Jantzer-Ward
Project Manager: Adriane Wineinger
Graphic Specialist: Karly Prickett

Advertising Sales

Technology Market Leader: Peg Miller
Key Account Director:
Chrissy Ferraro • 970-203-2883
Account Executives:
Megan Key • 970-203-2844
Barbara Ritter • 858-367-8058
Cass Schulz • 858-357-7649

Client Services

Senior Client Services Manager:
Michelle Andrews • 970-613-4964
Ad Production Coordinator: Kara Walby

Marketing & Circulation

Customer Service • 800-793-5697
Vice President, User Marketing &
Marketing Analytics: Tricia Syed
Marketing Director: Amy Connell

Technology Division & Penton Marketing Services

Senior Vice President: Sanjay Mutha

Corporate

Chief Executive Officer:
David Kieselstein
Chief Financial Officer/Executive Vice
President: Nicola Allais



List Rentals

MeritDirect
333 Westchester Avenue,
White Plains, NY 10604

Reprints

Reprint Sales:
Wright's Media • 877-652-5295

Windows IT Pro, November 2013, Issue No. 231,
ISSN 1552-3136. *Windows IT Pro* is published monthly by
Penton. Copyright ©2013 Penton. All rights reserved. No
part of this publication may be reproduced or distributed
in any way without the written consent of Penton.

Windows IT Pro, 748 Whalers Way, Fort Collins, CO 80525,
800-621-1544 or 970-663-4700. Customer Service:
800-793-5697.

We welcome your comments and suggestions about the
content of *Windows IT Pro*. We reserve the right to edit all
submissions. Letters should include your name and
address. Please direct all letters to letters@windowsitpro.com. IT pros interested in writing for *Windows IT Pro* can
submit articles to articles@windowsitpro.com.

Program Code: Unless otherwise noted, all programming
code in this issue is ©2013, Penton, all rights reserved.
These programs may not be reproduced or distributed
in any form without permission in writing from the
publisher. It is the reader's responsibility to ensure
procedures and techniques used from this publication are
accurate and appropriate for the user's installation. No
warranty is implied or expressed.

Windows®, Windows Vista®, and Windows Server®
are trademarks or registered trademarks of Microsoft
Corporation in the United States and/or other countries
and are used by Penton, under license from owner.
Windows IT Pro is an independent publication not
affiliated with Microsoft Corporation. Microsoft
Corporation is not responsible in any way for the editorial
policy or other contents of the publication.

Windows IT Pro

Surface 2 and Surface Pro 2



Paul Thurrott

is senior technical analyst for *Windows IT Pro*. He writes the SuperSite for Windows, a weekly editorial for *Windows IT Pro UPDATE*, and a daily Windows news and information newsletter called *WinInfo Daily UPDATE*.

Email



Twitter



Website



Microsoft's second-generation Surface tablets—the Windows RT 8.1-based Surface 2 and the Windows 8.1 Pro-based Surface Pro 2—were, at the time this article was written, to be made available for purchase on October 22, 2013, in many markets around the world. So what's new for this generation of devices? Here's what you need to know.

Worldwide Launch

Because Microsoft was new to the PC hardware market with the first-generation Surface devices, it launched those products in a measured way. Surface with Windows RT—Surface RT—launched first, in October 2012, and then only in the United States and Canada. Surface Pro didn't launch until February 2013. Over the intervening months, Microsoft has expanded the availability of its Surface devices geographically, of course, but also via other means, including the ISV channel.

With the second-generation Surface lineup, the availability picture has improved dramatically. Both Surface 2 and Surface Pro 2 will be launched together, for starters. But they're also launching in more markets simultaneously: 22 of them—including Australia, Belgium, Canada, Denmark, Hong Kong, Italy, and the United States—on October 22, followed by China in early November and elsewhere in the coming months. Let's look at the new models.

Still Two Models . . . For Now

Microsoft announced the two new Surface models, each of which replaces a previous Surface model, in September. Surface 2 is the second-generation version of Surface RT, and it utilizes an ARM processor and will run Windows RT 8.1. Surface Pro 2 is the second-generation version of Surface Pro; it features a fourth-generation Intel Core (“Haswell”) processor and will come with Windows 8.1 Pro.

Microsoft is also keeping its first-generation Surface RT device in the market as a low-cost option. (I'll discuss pricing later.) But I'm wondering whether that isn't temporary. Sometime before the end of the year, Microsoft will launch a third Surface model, which I think of as "Surface mini," which will also use an ARM (Qualcomm) processor and run Windows RT 8.1 and feature an 8" screen. It's possible—even likely—that this device will replace Surface RT as the new low-end model when it appears.

Shared Features and Technologies

Despite different hardware architectures, OSs, and target markets, Surface 2 and Surface Pro 2 share a surprising number of commonalities. Both devices still utilize the high-quality "VaporMg" casing made of magnesium, an integrated kickstand, front- and rear-facing cameras, multiple built-in sensors, and a 10.6" widescreen display.

But many of these common features have changed for generation 2. The VaporMg casing on Surface 2 is now a light gray color that Microsoft says is the natural color of the magnesium process. (Surface Pro 2 continues forward with the dark titanium color from before.) On both machines, the Windows logo has been replaced by a Surface logo: "Surface users want others to know they're using a Surface," Microsoft's Panos Panay said at the launch event.

That integrated kickstand has been updated, too: Instead of supporting just a single viewing angle, it now supports two. The second of the two viewing angles is geared toward a sorely missing capability on the first Surface devices: You can now use it on your lap. It's as solid as before and utilizes the same reassuring click.

Both devices sport USB 3.0 (in generation 1, Surface RT offered USB 2.0), Wi-Fi 802.11a/b/g/n, and Bluetooth wireless networking components. And the cameras in each machine have been updated. But the biggest change is the screen: Both devices now use a ClearType Full HD display, meaning 1080p or a resolution of 1920 × 1080. Surface Pro did offer this resolution before—Surface RT was a lowly 1366 × 768—but

the new screen offers 46 percent better color accuracy, so it's an improvement across the board.

Of course, Surface 2 and Surface Pro 2 do target different markets as well. So let's further examine where these devices diverge.

Surface 2

Surface 2 (see Figure 1) is Microsoft's answer to the Apple iPad, or what Panay called the "personal tablet" market. Of course, given how sales have gone over the past year, Microsoft is a bit sensitive about head-to-head comparisons of Surface with iPad, so the firm is also differentiating its product, once again, with Office. So, yes, Surface 2 is a personal tablet. But it's also the most productive personal tablet, one that ships with a free and full copy of Office, which features Word, Excel, PowerPoint, OneNote, and now Outlook. Microsoft appears to have addressed some of the original device's issues, and in ways that are fairly meaningful. Indeed, Panay said that Surface 2 was not a "subtle" upgrade from Surface RT but rather a complete revamp.

That revamp starts with the aforementioned 1080p screen, which makes a surprisingly strong difference, whether you're looking at text, graphics, or video. The Surface 2 also includes a much more powerful

Figure 1
Surface 2



TEGRA 4 processor that Panay says results in “unprecedented” performance for a personal tablet. “This is the fastest personal tablet [there is],” he said. “There’s no lag and no latency.”

Surface 2 also provides 25 percent better battery life than its predecessor, which Panay says is 12-plus hours in the real world. (Microsoft’s more conservative marketing materials state 10 hours as the official figure.) And the device itself is thinner and lighter than Surface RT, too.

Thinner, lighter, faster, and a much better screen: And Surface 2 comes with Windows RT 8.1, which is a significant improvement for both personal tablet users—who want to stay out of the desktop more easily—and traditional PC users alike. But it’s still not clear why Microsoft continues to push a product that, so far at least, hasn’t resonated in the slightest with consumers.

Panay mentioned that the app count in the Windows Store has grown from 10,000 at launch a year ago to 100,000 this year. But Windows “Metro” apps—which are the only app type you can install on Windows RT—haven’t taken off like Android or iOS did, and most big-name apps and games are still missing in action. The value proposition of this platform, as with last year, is still questionable, especially for consumers. For business users, Surface 2 is like a smartphone or, yes, a personal tablet, in that it’s managed by using Exchange Active-Sync (EAS) or MDM solutions such as Windows Intune. That might not be of interest, especially in larger organizations. Which is why Microsoft makes a Pro version of the device as well.

Surface Pro 2

Positioned as a new kind of ultrabook, Surface Pro 2 is less different from its predecessor than is Surface 2 compared to Surface RT. It utilizes the exact same body as the first Surface Pro, which is initially disappointing given its thickness and heft but makes sense when you consider the accessory compatibility story (especially with the new docking station).

Panay described the Surface Pro 2 (see Figure 2) as a modest upgrade by design: The original Surface Pro, he said, was “the best product Microsoft has ever made,” which is defensible, and “the best-selling device in its class,” which raises all kinds of questions. The Surface Pro was “working,” he said, and people understand it. So Microsoft put its energy into improving it, not reinventing the wheel.

Figure 2

Surface Pro 2 (with
Purple Type Cover 2)



To that end, the updates in this device are predictable. It features a fourth-generation Intel Core i5 “Haswell” processor, which provides a 20 percent performance boost but more impressive gains in battery life: 75 percent more, in fact, which means the device should get seven hours of battery life in real-world use. No real surprises, though as you’ll see the lineup has expanded to four models, some terrifically expensive.

Aimed at the business and “pro” markets, Surface Pro 2 is a real PC. This means that the device includes everything that’s good about PCs—100 percent compatibility with all Windows-compatible applications, utilities, and hardware peripherals; well-understood and fine-grained

management; full enterprise compatibility—and all of the bad, which includes viruses and other malware, and, of course, complexity.

But whereas a 10.6" screen makes some sense on a personal tablet, it makes little sense for this market. Today's mainstream ultrabooks are 13" or bigger, and available 11" devices, while slightly larger than the Surface devices, look like postage stamps by comparison. Some will appreciate the portability of such a device. But the combination of such a high-res but tiny screen and the Windows desktop, which is still ill-suited to such usage, is painful on all but the sharpest eyes. It just doesn't work.

Pricing

I had expected Microsoft to rejigger its Surface pricing structure in the wake of the past year's sales debacle, but the company did no such thing. Surface 2 pricing starts at \$449 for a 32GB model, just \$50 less than the Surface RT launch price and a full \$100 *more* than the price of that device today. A 64GB model is \$549. Both of these devices feature 2GB of RAM, which is just fine for Windows RT.

Surface Pro 2 is even worse. Here, Microsoft didn't lower the price at all, compared to the Surface Pro launch prices. A base 64GB version (with 4GB of RAM) costs \$899, while a 128GB version (also with 4GB of RAM) sits right at \$999.

Mind you, these prices don't include a typing cover of any kind, so the cheapest real-world price for this device is really \$1,028. You can get an 11" Macbook Air for \$999, and that device features a nice hardware keyboard and twice the storage. See the problem?

Responding to a need I'm not sure exists, Microsoft is also offering two higher-end versions of the Surface Pro. These versions bump the RAM to 8GB but oddly offer no improvement to the stock i5 processor.

A version with 256GB of storage will retail for \$1,299. And a limited-availability 512GB version will cost a whopping \$1,799, a price that would make even Apple blush.

Accessories

Microsoft also unleashed new Surface accessories, most of which work with the original generation devices. The two most eagerly awaited ones won't ship until next year: A new docking station (\$199) for Surface Pro and Pro 2 (see Figure 3) offers a USB 3.0 port, three USB 2.0 ports, audio in and out, and a mini DisplayPort 1.2 connector that supports a 3840 × 2160 external screen. A much-needed Power Cover (\$199), essentially a Type Cover with a battery, is compatible with Surface 2, Surface Pro, and Surface Pro 2 (but not Surface RT) and will provide two to four hours of additional battery life, depending on device.

Figure 3
Surface Pro Docking
Station



Several other accessories will be available immediately or in October. Touch Cover 2 is thinner than its predecessor but offers automatic backlit keys and will cost \$119. It comes only in black and works with all Surface models. Type Cover 2 is also thinner than its predecessor and features automatic backlit keys. It comes in a new felt-like exterior that covers the front and rear of the cover. Available in magenta, cyan, purple, or black, this accessory will cost \$129. It works with all Surface models.

Microsoft is also providing a tubular wireless adapter for Type Covers, which lets you use a Type Cover, Type Cover 2, Touch Cover,

Touch Cover 2, or Power Cover detached from the Surface device at a distance of up to 30 feet. It works with all Surface models and costs \$59. And a new car charger with USB (\$49) extends the lineup of Surface chargers; it also works with all Surface models.

Early Advice

Although I'll be reviewing Surface 2, Surface Pro 2, and various accessories in the coming weeks, my hands-on time with the devices this week is enough to offer some basic advice. If you already own a first-generation device, there's no reason to upgrade, though a Power Cover (and possibly a docking station) will be a welcome addition for any Surface Pro user.

Although I'm interested in Surface 2, I recommend that potential customers wait and see what the Surface "mini" looks like: This smaller form factor might make more sense for a personal tablet, depending on your needs. And Surface Pro 2? I don't know. Although everyone's needs are different, I think the 10.6" screen is just too small for regular productivity use, and I eagerly await a 13" device—but with the understanding that it probably won't happen anytime soon, if ever. These devices are all incredibly well made and feature innovative designs. I'm just not sure that a market exists for either as they now stand. ■

PowerShell Cmdlets for DNS

Three tools (out of 100!) that simplify your DNS administration



Mark Minasi

is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books, including *Mastering Windows Server 2008 R2* (Sybex). He writes and speaks around the world about Windows networking.

Email



Twitter



Website



If you perform DNS administration on Windows Server systems, you might remember the columns I wrote years ago—such as “[Scripting DNS Setup](#)”—about Dnscmd, a powerful command-line tool that lets you do just about any kind of DNS administration. Now, I don’t want to beat up on Dnscmd, but I was pretty happy to see that my command-line DNS administration work now has a new assistant, and that assistant comes in the form of 100 PowerShell cmdlets. This month, I want to show you 3 of those 100 cmdlets, as well as a couple examples of how you can take a simple PowerShell “power tool” and make it even more powerful.

Set ‘Em Up

First things first. What version of Windows Server will you need in order to use these 100 new cmdlets? I haven’t yet tried all 100 of them, but as far as I can see, all you need to run these DNS-oriented cmdlets against a Windows Server 2008 R2 system is a [Windows 8](#) workstation that you’ve joined to a domain and upon which you’ve downloaded and installed the Remote Server Administration Tools (RSAT) for Windows 8. At that point, type:

```
get-command *-dnsserver*
```

and you’ll see the 100 DNS cmdlets. Of course, if you already have a [Windows Server 2012](#) system, just install the built-in DNS management tools and you’ll get the DNS cmdlets. OK, let’s put them to work.

Two Gets

By now, you've probably seen that a good strategy for testing a block of PowerShell cmdlets is to give the Get- cmdlets a try. Get-DNSServerEdns tells you whether your DNS server has *Extensions to DNS* enabled (and it should, as in every case I've come across). Get-DNSServerRecursion tells you whether your DNS server will accept queries for zones that don't exist on that server.

Suppose I decided to ask my local AD-centric DNS server to find www.microsoft.com. I don't run Microsoft's public zone, so unless I have recursion enabled on my local DNS server, the server won't go running around the Internet to resolve my query.

Why would this functionality be useful? Unfortunately, the number of jerks out there using DNS servers to create a nasty kind of distributed, "amplified" Denial of Service (DoS) attack has grown tremendously in the past few years, prompting many of us to have to disable recursion (it's on by default) on our DNS servers and find resolvers elsewhere. In such a case, it's nice to be able to check on a given DNS server or, of course, to use PowerShell's great remoting tools. Thus, if you had, say, five DNS servers named D1, D2, D3, D4, and D5, you could get one simple report of their recursion status with this command:

```
invoke-command -computername d1,d2,d3,d4,d5 -scriptblock
    {get-dnsserverrecursion|select pscomputername,enable}
```

I've shown you the *invoke-command* cmdlet before, and as I said then, the *-scriptblock* parameter tells the remote computers what command to execute between the braces. In those braces, you see the DNS cmdlet and then a Select statement that shows you the name of the computer that you ran this cmdlet on, as well as the result.

One Add

I found another Server 2012 DNS cmdlet, *add-dnssecondaryzone*, to be quite a time-saver recently. I was setting up a new DNS server that

I wanted to configure as a secondary server for a bunch of the DNS zones that I host.

Don't misunderstand: Making a Windows DNS server into a secondary server for an existing domain isn't difficult to do with the GUI or `Dnscmd`, but it's tedious. It took just a moment to figure out that I could make the DNS server that I was sitting at a secondary DNS server for a domain named `bigfirm.com` with a primary IP address of `71.23.1.5`, and that I wanted its zone data stored on a text file named `bigfirm.com.dns`:

```
add-dnssecondaryzone bigfirm.com "bigfirm.com.dns" 71.23.1.5
```

I immediately thought, OK, I have to do that with seven domains, and all I have to do is change that zone name in two places. So, let's see, first I'll store the domain names in an array, which is nothing but a list of the domain names in quotes with commas between them:

```
$zones="bigfirm.com","minasi.com","mmco.com"  
      "pungogrill.com","thesoftwareconspiracy.com",  
      "softwareconspiracy.com","steadierstate.com"
```

and then I can use the pipeline and the *foreach-object* command to feed each zone name to `add-dnssecondaryzone`, using the built-in `$_` variable that contains whatever is in the pipeline at the moment. That lets me type this:

```
$zones | foreach-object {Add-DnsServerSecondaryZone $_ $_+".dns"  
                        "71.23.1.5"}
```

I freely admit that I got a bit ahead of myself there, but believe me, it worked, and on the first try! So I think it's about time to get off my duff and cover `ForEach-Object` and the `$_` pipeline tools—next time. ■

Top 10 Ways Microsoft Can Fix the Surface Pro

Surface Pro 2 is already approaching, but these are the top-of-mind changes I'd like to see

I am a long-time gadget-head, but I must admit I've been a bit put off by today's crop of tablet devices. All these devices are pretty good for consuming content, browsing the web, playing games, or doing light email work, but what if you need to create content and write .NET code? Today, only the Surface Pro can really accommodate my requirements. I've owned one for a few months now, and [I'm finding a lot to love about the Surface Pro—but a lot that frustrates me](#). Microsoft is getting ready to release a new [Surface Pro 2](#), but if I were in charge of the company, here are the top 10 things I would do to fix the Surface Pro.



Michael Otey

is senior technical director for *Windows IT Pro* and *SQL Server Pro*.

① Make It Thinner and Lighter

Although it's thinner and lighter than most laptops, the Surface Pro is noticeably thicker and weightier than either the Apple iPad or the Surface RT. It's certainly more powerful than either of those tablets, but it really needs to match the portability factor. Weight might be the main selling point of these types of tablet devices.

② Give It a Bigger Screen

Yes, I know there's a trend toward smaller iPads, and there are rumors of a smaller-form-factor Surface on the horizon. But that's not what I want. After working with the Surface Pro for the past few months, I often find that I want slightly *more* screen real estate, not less. But remember that I create content; a larger screen helps me to be more productive. I'd like to see something closer to a full 8" × 11" form factor.



Email

③ Beef Up the Screen Resolution

While I'm on the subject of screens, the Surface Pro isn't shabby in this regard, providing high-definition 1080p (1920 × 1080) resolution, but it competes with the iPad 4 and Nexus 10, which provide 2048 × 1536 and 2560 × 1600 resolutions. No one is really sold on how ClearType evens this out. The Surface Pro is a more expensive device, and frankly it needs to compete better in the screen-resolution numbers game.

④ Provide More Disk Capacity

I definitely love the SSD in the Surface Pro. It's high-performance, and it boots and resumes very quickly. However, I do wish it were bigger than 128GB, especially considering the fact that after the OS and the recovery partition are installed, there's only about 83GB of useable storage left. A 250GB SSD option (or larger) would help solve this problem.

⑤ Give It More Battery Life

I have no complaints about the performance of the Surface Pro. However, there are times when I wish it offered more battery life. Currently, the unit lasts about as long as a laptop—not as long as an iPad. The more power-efficient Surface RT lasts quite a bit longer than the Surface Pro. The Surface Pro needs either a more power-efficient processor or a higher-capacity battery.

⑥ Put an Insert Key on the Keyboard

I really like the magnetic clip-on Touch and Type keyboards, but why did Microsoft omit a key? The company can get away with this omission on the Surface RT, which doesn't have backward compatibility, but many existing applications actually use the Insert key. While I'm on the subject of the keyboard, the optional Type cover could really be sturdier—I broke some of the keys off mine just by carrying it around by itself in my computer bag.

7 Give Us at Least One More USB Port

Although this is an area where the Surface beats the port-less iPad hands-down, it could still do better. One port is the absolute minimum. Most Surface users I know carry around a USB hub to help address this shortage. The bottom line is that you wind up wanting a USB port on each side, and more than one would be better.

8 Make the Magnetic Charge Port Wider

The Apple magnetic charge port is brain-dead simple to use. Whenever you get the charger close to the port, it simply snaps into place. Not so with the Surface Pro. It's difficult to get that narrow attachment in correctly. This might seem like a small point, but the Surface competes with the iPad, and because it's number two (or three)—and more expensive—it needs to do better, or at least work well, in every area that users will notice. And they will notice.

9 Offer More Apps

Perhaps this is a bigger weakness with Windows Phone, but it's a weak point for the Surface and Windows 8 in general. Apple and Android devices have access to 800,000 or more apps. Optimistic estimates put Windows Store apps at about 160,000. That's not a tiny number, but it's nowhere near what the other tablets offer.

10 Revamp the Windows 8 Interface

Maybe Microsoft should have watched that Paul Masson commercial, in which Orson Welles said, "We will sell no wine before its time." It seems to me that Microsoft released Windows 8 before it was ready. The hybrid touch/desktop experience is difficult for users to grasp, and Microsoft didn't help them by jettisoning the familiar Start button. That said, Microsoft has shown an ability to rectify these types of shortcomings. Windows 7 was a terrific replacement for the widely reviled Windows Vista release. Perhaps Microsoft can repeat this magic with Windows 9. ■

The New Microsoft and the Future of the Identity Professional

What does Microsoft's new strategic direction mean for the "ID pro"?



Sean Deuby

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel's core Directory Services team. He's been a Directory Services MVP since 2004.

Email



Twitter



In the past few months, [Microsoft has made a number of sweeping changes](#) that will have long-term consequences for the company, the businesses that depend on Microsoft, and the IT professionals that support Microsoft's products. How will these changes affect identity professionals in the Windows world?

Shifting Tides

Microsoft has made it quite clear that it is evolving from an "on-premises software company" to a "devices and services company." The company has been beefing up its hardware product line, from its keyboards and mouse devices to its tablet/PC devices; my colleague Rod Trent believes that [Microsoft will begin to build servers as well](#). It has acquired Nokia's devices and services business, putting Microsoft squarely in the smartphone market as number three after Google and Apple. And it has eliminated a couple of important programs for IT pros: [TechNet software subscriptions](#) and the [Microsoft Certified Master program](#). I'm sure these aren't the last of the strategic changes we'll see coming out of Redmond this year.

To gauge how these changes and the overall impact of services versus traditional on-premises systems will affect the identity professional, let's compare identity to another infrastructure system: messaging. Running a robust and reliable email infrastructure isn't a trivial or

inexpensive task, and letting someone else handle the complications is very appealing. In addition, messaging has the Achilles heel of spam filtering, spear phishing, and Denial of Service (DoS) attacks to deal with; protection against these vectors is ever-changing and can get very expensive to manage yourself. As a result, hosted email has become one of the most popular cloud service migration targets.

As with email, the future of the corporate identity professional's job role depends on the on-premises/cloud equation. The difference between email and identity however, is that—unlike email—identity is a critical part of the entire IT infrastructure, not just one aspect of it. Because Active Directory (AD) is a major part of on-premises identity, let's focus on how on-premises/cloud impacts it.

Active Directory Impact

With an on-premises capability, you must handle it all. You must deploy, support, and upgrade the service infrastructure (which means capital hardware and software costs), then pay for sustaining software maintenance and human resource costs. You must also manage the service's data—archiving, recovery, and governance. For AD, this means maintaining both the service (the network of domain controllers—DCs) and the logical administrative structure of AD (the identity and network data it contains). In the largest enterprises, the management of the AD service itself and the data within it are delegated to two or more groups (e.g., AD infrastructure support and account management). How do cloud identity services affect these jobs?

In the simplest terms, as long as you've established on-premises applications, you'll want to have an on-premises identity infrastructure. (It's possible to have on-premises applications with completely outsourced identity, but I see that as an edge case.) The size of your identity infrastructure might decrease as you move your computing to an external provider, but it won't go away; there are too many perfectly functional production applications that show little to no return on investment (ROI) for moving off premises, and they'll need identity

to work. Even if you could easily move production virtual machines (VMs) from your data center to the cloud (and it will become much easier as [Windows Server 2012](#) plus Windows Azure IaaS gains traction in the enterprise), why would you want to start paying for Windows Azure when you've already invested in your own data center and have depreciated hardware? Over time, this scenario might become more inviting as your data center hardware approaches obsolescence—but not immediately.

Keys in the Cloud

An equally important factor is security. Do you want to store the keys to the kingdom (your company's identities and passwords) in the public cloud? What is far more likely is that your cloud computing will be of the “private cloud” type for the near future. And the identity infrastructure required for an on-premises private cloud is every bit as complicated, if not more complicated, than what you have today.

If you build a hybrid cloud strategy, you might well have an identity presence in the cloud (e.g., Windows Azure AD or another IDaaS provider) and therefore some identity data management responsibilities without the corresponding service management. But with that scenario will come identity-bridge (the connection between local AD and cloud identity services) management responsibilities.

The AD data-management aspect is a little less exciting, which in this context about career security is probably a relief. Whether you're on premises or in the cloud (or both), you still have the account management lifecycle—also known as “create, read, update, delete” (CRUD)—to deal with. It's simply the mechanisms to handle this lifecycle that will change.

SSO to Cloud


Finally, the recent revelations about NSA surveillance have only underscored the importance of using federation to provide single sign-on (SSO) to cloud service providers. A federated trust between

identity provider and relying party does not provide passwords to the relying party (aka a service provider such as Tripit.com). This is why I don't recommend the [password synchronization option](#) of Microsoft's Dirsync account synchronization utility for [Office 365](#) or other Windows Azure-based services; you should instead use Dirsync in combination with Active Directory Federation Services (AD FS) or another federation service to create a federated trust between you and Windows Azure AD.

I'm not worried about the future of the identity professional. In fact, the role of identity and the need to comprehend it well has only grown in the past few years. The installed base of AD might shrink somewhat over the coming years, but it'll be a long time before we talk about AD as one of those "used to have" technologies. ■

CAN'T GET AWAY?

Get first-class education from your desk



Windows IT Pro offers FREE online events including webcasts, demos and virtual conferences. All events are brought to your computer live while being fully interactive.

Go to www.windowsitpro.com/events to see an up-to-date list of all online events.

Windows IT Pro

First-class education from the top experts in the industry.
Visit www.windowsitpro.com/events for a knowledge upgrade today!

Have a full plate on the live date? Don't sweat it! All online events are recorded and available 24/7.

Failover Clustering in Windows Server 2012 R2

New features and enhancements abound



**John
Marlin**

is a senior support escalation engineer in Windows Commercial Technical Support, focusing on failover clustering. He is a Microsoft Certified Trainer for clustering, delivering to Microsoft and its partners, and is a regular contributor to the [Ask the Core Team](#) blog. He is also a contributor to the new book *Introducing Windows Server 2012* (Microsoft Press).

Email



Website



Windows Server 2012 R2's new features and enhancements for failover clustering were made with easier management, increased scalability, and more flexibility in mind. Here are the most noteworthy changes.

New Shared .vhdx Files

One new feature that seems to be getting the biggest raves is the ability to have shared Virtual Hard Disk (VHD) files (in the .vhdx file format) for guest clusters running in Hyper-V host clusters. What this means is that you can now use VHDs for your guest clusters without having to also attach your actual storage to those virtual machines (VMs). The shared .vhdx files must reside on the local Cluster Shared Volumes (CSVs) of the host cluster or on a remote Scale-Out File Server.

When you create a .vhdx file for a VM, there's a new option to mark it as shared. If you're using Hyper-V Manager, you select the *Enable virtual hard disk sharing* option in Advanced Features in the VM's settings, as Figure 1 shows. If you're using [Microsoft System Center](#) Virtual Machine Manager (VMM), you select the *Share the disk across the service tier* option on the Hardware Configuration page, as Figure 2 shows. You then add the same .vhdx file to each of the other VMs and select the same setting. When you attach these shared VHDs to guest VMs, they'll appear as Serial Attached SCSI (SAS) drives to the guest VMs.

You can use [Windows PowerShell](#) to set up the .vhdx files if desired. For example, suppose that you want to create a 30GB .vhdx file and assign it to two VMs as a shared VHD. First, you need to create the .vhdx file by running a command like this:

```
New-VHD -Path C:\ClusterStorage\Volume1\Shared.VHDX `
-Fixed -SizeBytes 30GB
```

Then, to assign it to each VM as a shared .vhdx file, you use commands like these:

```
Add-VMHardDiskDrive -VMName Node1 `
-Path C:\ClusterStorage\Volume1\Shared.VHDX `
-ShareVirtualDisk
```

```
Add-VMHardDiskDrive -VMName Node2 `
-Path C:\ClusterStorage\Volume1\Shared.VHDX `
-ShareVirtualDisk
```

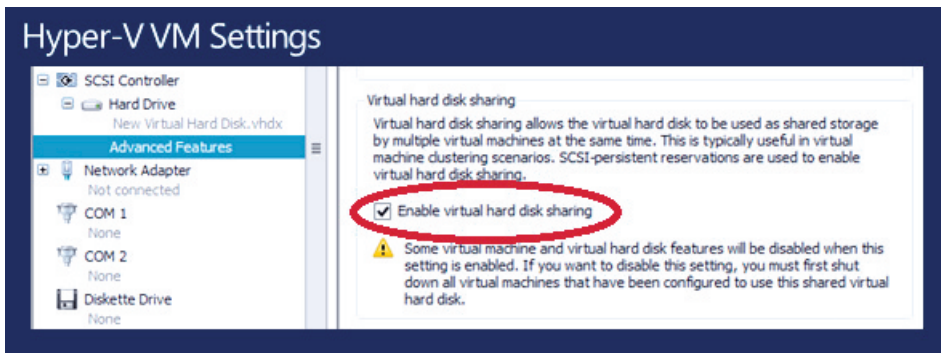


Figure 1

Enabling .vhdx File Sharing in Hyper-V Manager

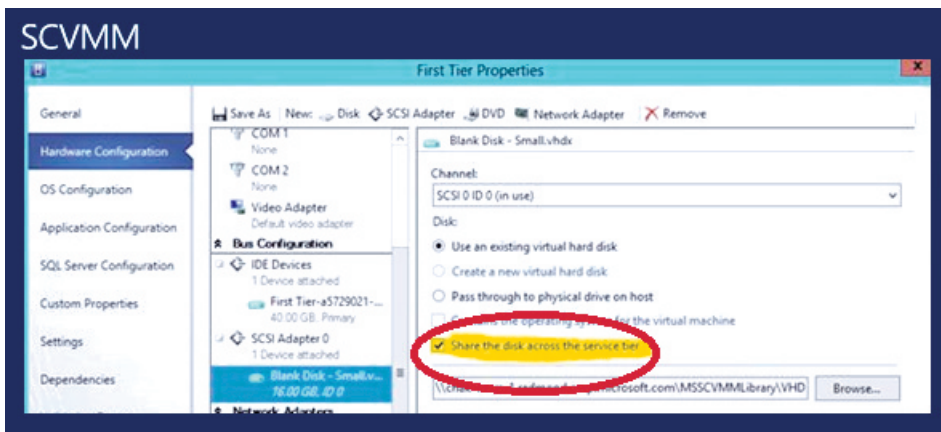


Figure 2

Enabling .vhdx File Sharing in VMM

Using shared .vhdx files is ideal for:

- File services running within a VM
- SQL Server databases
- Other database files that reside on guest clusters

You can find more information about shared .vhdx file requirements and configurations in the [Virtual Hard Disk Sharing Overview](#) web page.

The new Clusters dashboard makes it easier to manage multi-cluster environments.

New Node Shutdown Process

When you use failover clustering in [Windows Server 2012](#) and earlier, Microsoft recommends that you first move all the VMs off a node before you shut it down (or reboot it). Here's why: When you shut down a node, it induces the cluster-controlled action of Quick Migration on each VM. A quick migration will put that VM in a saved state, move it to another node, then make it come out of the saved state.

When a VM is in a saved state, it's actually down, which means your productivity is down until the VM comes back online. The reasoning behind the recommendation of moving all VMs off the node before shutting it down is that you can use live migration to move those VMs so that there's no loss in productivity. However, if you follow this recommendation, shutting down a node can be a lengthy manual process.

In Server 2012 R2 failover clustering, Microsoft has changed what happens when a node is shut down. The new process has two main components: drain on shutdown and "best available node" placement.

If you shut down a node without first putting it into maintenance mode, the cluster will then automatically issue a drain. During the drain, the cluster uses live migration to move the VMs off the node, following the machine priority (high, medium, and low). All the VMs are moved, including the low-priority VMs.

When migrating the VMs, the cluster uses "best available node" placement. Here's how it works: Before the cluster starts migrating the VMs, it first checks the available memory of the remaining nodes. Using this information, it strategically places the VMs on the best




Figure 3
Migrating the VMs
to the Best Available
Node

available node, as Figure 3 shows. This ensures a smoother transition because it prevents moving high-priority VMs to a node that doesn't have enough memory.

This new process is enabled by default. If you need to manually enable or disable it, you can configure the `DrainOnShutdown` cluster common property. To enable it, use the PowerShell command:

```
(Get-Cluster).DrainOnShutdown = 1
```

To disable it, run the command:

```
(Get-Cluster).DrainOnShutdown = 0
```


Additional Health Detection Feature for VM Networks

With Server 2012 R2 failover clustering, you have an additional health detection feature for the networks the VMs utilize. If a node's network goes down, the cluster will first check to see if the network is down

across all the nodes. If it is, the node's VMs will remain where they are. If this is the only node with the problem, the cluster will use live migration to move its VMs to a node in which the network is available.

This feature is enabled by default on all networks configured for VMs. If there are any networks that you don't want to protect with this feature, you can disable it using Hyper-V Manager. You simply need to clear the *Protected network* check box in Advanced Features in the VM's settings, as shown in Figure 4.

Figure 4
Disabling the
Protected Network
Feature



New Clusters Dashboard

When managing multiple clusters in Server 2012 and earlier, you must switch between each cluster to see whether there are errors or any concerns. That's no longer the case in Server 2012 R2 because Failover Cluster Manager has the new Clusters dashboard, which Figure 5 shows.

The new Clusters dashboard makes it easier to manage multi-cluster environments. You can quickly check the status of roles and nodes (e.g., up, down, failed) and see whether there are any recent events to review. Everything is hyperlinked, so simply clicking the link takes you to what you need to see. For example, clicking the *Critical: 3, Error: 1, Warning: 2* link shown in Figure 5 brings up the list of these filtered events so that you can go through them.

Figure 5
Using the New Clusters
Dashboard

Clusters			
Name	Role Status	Node Status	Event Status
2012R2-CLUSTER.CONTOSO.COM	0 total	1 total	Critical: 3, Error: 1, Warning: 2
TEXAS-CLUSTER.CONTOSO.COM	0 total	1 total	None in the last hour

CSV Enhancements

In Server 2012 R2 clustering, Microsoft has made several CSV enhancements. They include optimizing the CSV placement policy and adding a dependency check.

The CSV placement policy now spreads ownership of the CSV drives among the nodes to ensure they're evenly distributed. For example, suppose that you have three nodes with four CSV drives, each of which houses five VMs. When all the nodes are running, two of the nodes have one CSV drive and five VMs. The other node has two CSV drives, each with five VMs. You now need to add a fourth node to the cluster. As soon as you add the node, the cluster will automatically give this new node ownership of one of the CSV drives. All the VMs running on that CSV drive will then be moved to this new node using live migration. By doing this, the cluster has more evenly spread the load among all the nodes.

Another enhancement that Microsoft made to CSVs is adding a dependency check. When a node isn't the owner (or coordinator) of a CSV drive, it must go through the network with a Server Message Block (SMB) connection to the coordinator for any metadata updates needed for the drive. The coordinator node has an internal share to which all the other nodes connect for this purpose. This requires the Server Service to be running. If the Server Service were to go down for some reason, the non-coordinator nodes wouldn't have the SMB connection, which would cause errors. More important, any metadata updates would simply be cached rather than sent because there's no way to send them. To get out of this situation, you need to manually move ownership of the CSV drive to another node.

To help avoid this situation, Microsoft added a dependency check that monitors the health of both the internal share and the Server Service. If a dependency check reveals that the Server Service is down, the cluster will move ownership of any CSV drives that the node owns to other nodes. The cluster will also follow the optimized CSV placement policy to evenly distribute the CSV drives. For example,

suppose that you have a cluster with three nodes, each of which holds two CSV drives. If one of the node's Server Service goes down, the cluster will move ownership of that node's two CSV drives to each of the remaining two nodes.

Improvement in Network Validation Tests

Failover clustering has always used port 3343 for all communications (e.g., health checks, status reporting) between the nodes. However, there has never been a check for this port. The network validation tests only checked basic network connectivity between the nodes. Because these tests never checked for connectivity over port 3343, you wouldn't know if the Windows Firewall Port 3343 rule was disabled or that port 3343 wasn't open because of a third-party firewall being used.

In Server 2012 R2, the new Validation Network Connectivity test checks for communication over port 3343. When troubleshooting communication problems in the past, you might not have always checked this port first. With this test, it can be your first check. If the port is causing the problem, you'll have saved yourself quite a bit of time troubleshooting the problem.

Dynamic Quorum Enhancements

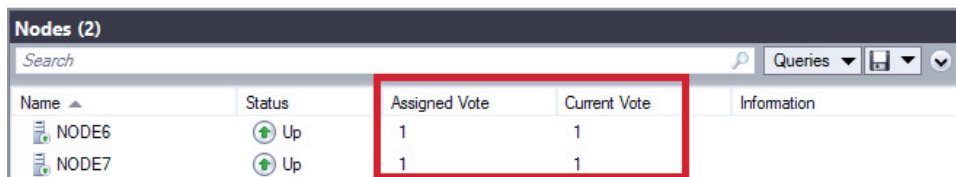
In Server 2012 failover clustering, Microsoft introduced the concept of the dynamic quorum. When the dynamic quorum feature is enabled, the cluster automatically adjusts the number of votes required to keep a cluster running if nodes go down. In Server 2012 R2 failover clustering, Microsoft has gone a step further by introducing the dynamic witness feature and the `LowerQuorumPriorityNodeID` property.

When the dynamic witness feature is enabled, the cluster dynamically adjusts the vote of the witness resource (a disk or file share). If there is a majority of nodes (i.e., an odd number of nodes), the witness resource will have its vote removed. If there isn't a majority of nodes (i.e., an even number of nodes) or if the witness resource is needed for a vote, it's dynamically given back the vote.

Because of the new dynamic witness feature, Microsoft has changed its witness recommendations. Previously, the recommendation was based on the number of nodes. If you had an even number of nodes, Microsoft recommended adding a witness resource to get to an odd number. If you had an odd number of nodes, it recommended not adding a witness resource.

With Server 2012 R2, the recommendation is to always add a witness resource. Because of the dynamic witness feature, the cluster will give the witness resource a vote if the cluster needs it or remove the vote if the cluster doesn't need it.

The cluster also adjusts the node weights as needed for when nodes go down or join the cluster. Because of the dynamically changing node weights, they've been added to Failover Cluster Manager so that you can quickly see these weights without having to run any commands against the nodes. You can see these values by selecting Nodes in Failover Cluster Manager, as Figure 6 shows. Note that you still have the option within the quorum configuration to remove a node's vote if desired.



Name	Status	Assigned Vote	Current Vote	Information
NODE6	Up	1	1	
NODE7	Up	1	1	

Figure 6

Checking the Node Weights

Another dynamic quorum enhancement has been made in the area of multi-site clusters. When you have nodes in two different sites and there's a network break between the two sites, only one site is going to remain running. In Server 2012 (and earlier) failover clustering, the site containing the node that gets the witness resource first is the site that remains running. However, this site might not be the one that you want to remain running. In other words, when you have a 50-50 split where neither site has quorum, you have no way of preselecting which site should remain running.

In Server 2012 R2 failover clustering, there's a new cluster common property that you can use to determine which site survives. You can set the `LowerQuorumPriorityNodeID` property to specify which node will have its vote removed in case of a 50-50 split.

For example, suppose you have three nodes in your primary site and another three nodes in an offsite location. You can set the `LowerQuorumPriorityNodeID` property on the offsite nodes so that if you have a 50-50 split, the offsite nodes will stop their Cluster Service until network connectivity is restored. To set this up, you first need to know the Node IDs of the offsite nodes. You can find out this information by running the following PowerShell command for each offsite node (where *NodeName* is the name of that node):

```
(Get-ClusterNode -Name "NodeName").Id
```

After running these commands, let's say that you find out the Node IDs of the offsite nodes are 4, 5, and 6. To ensure that these offsite nodes go down if you have a 50-50 split, you run these commands:

```
(Get-Cluster).LowerQuorumPriorityNodeID = 4  
(Get-Cluster).LowerQuorumPriorityNodeID = 5  
(Get-Cluster).LowerQuorumPriorityNodeID = 6
```

Now if you have a break in communication, the offsite nodes will stop their Cluster Service and all roles in the cluster will stay in the primary site nodes, which will remain running.

Even More Changes

Many new features and enhancements have been added to failover clustering in Server 2012 R2, and they're all for the better. I've introduced you to only some of them. If you want to learn about the changes I didn't discuss or want more information on those I've covered, see the [What's New in Failover Clustering in Windows Server 2012 R2](#) web page. ■

Windows Server 2012

Remote Server Management

Run graphical remote management tools on a client desktop

Windows Server 2012 introduced a staggering number of changes in virtualization features: an updated Hyper-V, desktop virtualization with new Remote Desktop Services, and much more. Apart from features, though, Microsoft made a big shift in the best practice deployment of Server 2012.

Since Windows Server 2008, you could use the Server Core installation option (which I cover in detail in “[Windows Server 2012 Installation Options](#)”) to install Windows Server without a local graphical interface; however, this was really only for certain infrastructure roles. Server 2012 changes this. Server Core is the default installation option for Server 2012—and is for more than just infrastructure roles.

Server Core is a target for application workloads such as SQL Server in Server 2012. You can add or remove the graphical interface and management infrastructure and tools at any time. The upshot is that in Server 2012, many more servers will be Server Core rather than Server with a GUI (the new name for what was Full Installation). In short, you probably should change the way you manage your Windows environment.



John Savill

is a Windows technical specialist, an 11-time MVP, and an MCSE for Private Cloud and Server Infrastructure 2012. He's a senior contributing editor to *Windows IT Pro* and his latest book is *Microsoft Virtualization Secrets* (Wiley).



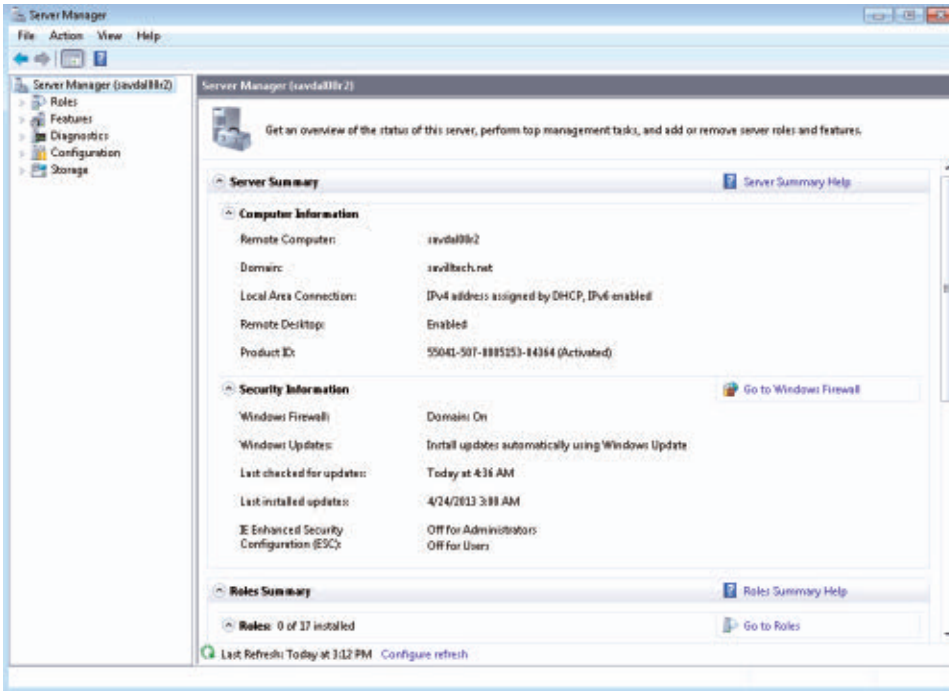
Evolution of Graphical Management

To put it into perspective, I look back at my first job working with computers. I was a VAX/VMS systems administrator, which really meant I changed backup tapes and provided printouts to developers. I managed the servers by sitting in the server room, where a video terminal was connected to each server. Then we added TCP and I would sit at a desk and telnet directly into each server. I managed Windows servers (when we got them) the same way. I had a monitor and keyboard in the server room that I would use to connect to a server via a KVM. And then when RDP and remote administration became a key feature, I worked from my desk and just RDP'd into the server. In all cases, though, I effectively managed one server at a time—and managed it locally.

As the number of server instances explodes, managing one server at a time becomes unrealistic. And with the shift to Server Core, local management with a graphical interface is no longer an option. We must shift how we manage servers.

Remote management (and by that I don't mean RDPing into the server, but rather running management tools on your client desktop that connect to the server) has been a *realistic* option since Windows Server 2008 R2. I emphasize *realistic* because it was *possible* before Server 2008 R2; however, it was cumbersome.

In Server 2008 R2, Server Manager finally allowed you to remotely connect to a server and manage it (providing the server has been enabled for remote management). This means you can install the Windows 7 Remote Server Administration Tools (RSAT) on your Windows 7 desktop, fire up Server Manager, and remotely manage a Server 2008 R2 operating system. However, this provides a connection to just one remote server at a time (Figure 1), which isn't convenient for managing multiple servers. Also, you might need to consider other types of applications that may also require management tools to run locally on the application server and not support Server Core; however, most applications have moved to a remote-management tool model or completely web-based administration.

**Figure 1**

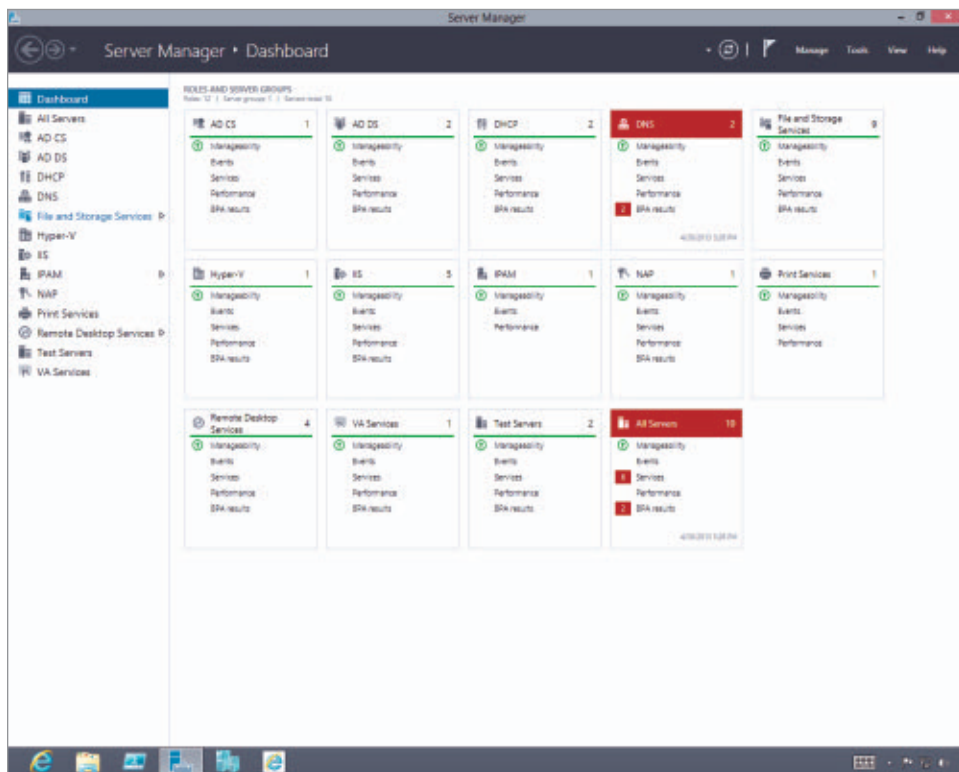
Using Server Manager to Manage Windows Server 2008 R2 from a Windows 7 Client

Remote management tools are important for another reason. New graphical management tools are very rich, but the reality is you typically wouldn't want them running on your servers because the tools would take up resources that would be better used by the services running on the OSs. It's far better to have that rich graphical interface rendering on your desktop OS and minimal network traffic in the process of fetching information and performing tasks.

Graphical, Multi-Server Management

Server 2012 embraces graphical, multi-server management with its completely reenvisioned Server Manager environment. A key reason for Server Manager's complete redesign was to accommodate the simultaneous management of multiple servers. Figure 2 shows Server Manager running on a Windows 8 machine with Windows 8 Remote Server Administration Tools (RSAT) installed. This is how Microsoft wants you to manage Server 2012—and how *you* should want to manage it.

Figure 2
Server Manager
Running on Windows 8
and Managing Multiple
Windows Servers



Manage multiple servers. If you consider that all your servers are running Server Core, then managing multiple servers with the ease of one should sound pretty good. In Figure 2, the initial dashboard not only shows server health, but Server Manager also automatically detects the roles installed on the servers it's managing and creates separate groups so you can see the health of Microsoft IIS, Hyper-V, and Active Directory (AD). You can even create your own groups of servers; for example, "Test Servers" in my instance. It's possible to click any problem areas (identified with red highlighting in Figure 2) and quickly see details. For example, click the Best Practice Analyzer (BPA) items in red to see what the problems are (yes, BPA is integrated with Server Manager) or which services are experiencing problems. All those services across the various servers can then be restarted with a single click.

There is one difference between Server Manager running on a Windows 8 client and Server Manager running locally on Server 2012 (Figure 3): the inclusion of a Local Server navigation node on Server 2012. Through this Local Server workspace, you can perform initial configurations on items such as server name and IP address (by *initial* I mean items you configure once on a server after installation and typically never change again). For this reason, items on the Local Server workspace aren't available via Server Manager's remote management functionality because they weren't deemed necessary for ongoing management. I expect that will change in the future. The good news is that remote management is enabled on Server 2012 by default; no manual action is necessary. This doesn't mean you can RDP into the server by default, but you can fire up Server Manager and (providing you have permissions on the server) remotely manage it.

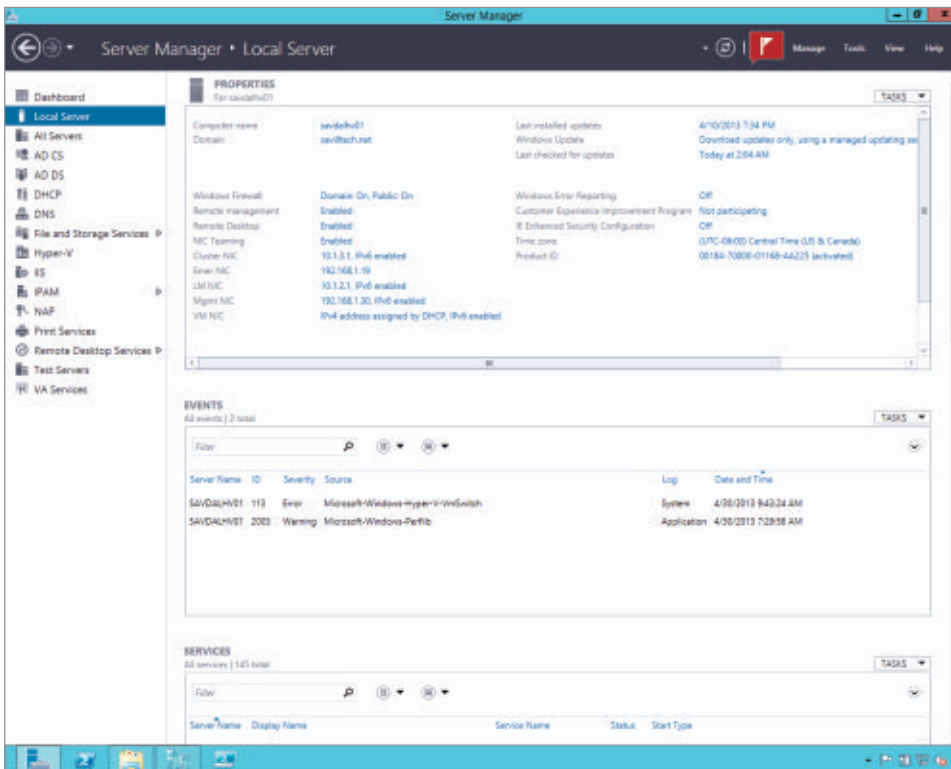


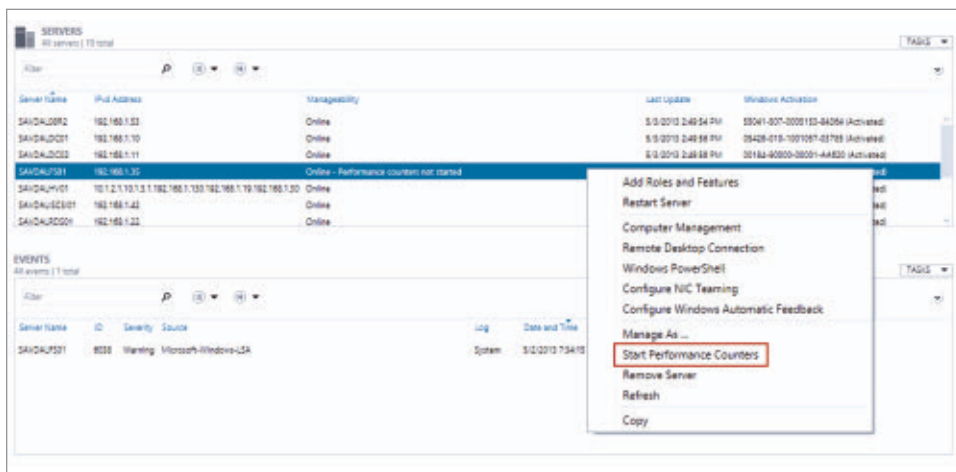
Figure 3
Server Manager
Running Locally on
Windows Server 2012

You can add servers to Server Manager via the Add Servers action on the Manage menu. While the focus of this article is remote management of Server 2012, you also can manage Server 2008 and Server 2008 R2 with Server Manager. The only requirement is Windows Management Framework (WMF) 3.0 must be installed on the Server 2008 and Server 2008 R2 boxes. You can [download Windows Management Framework 3.0](#) online, but make sure you've [installed .NET Framework 4.0](#) first, or WMF 3.0 won't install correctly.

As you add new servers to Server Manager, it polls them, detects the various infrastructure roles, and surfaces the roles in the appropriate navigation areas of Server Manager. Servers are exposed through the All Servers workspaces. It's also possible to create your own groups of servers, like the "Test Servers" group in Figure 3.

Make information visible. One action you'll likely need to take is right-click the server(s) and select Start Performance Counters (Figure 4), which enables a collection of performance data about the source server. You can access the data remotely through Server Manager to view key performance metrics such as CPU and memory use. You can even set Server Manager to alert you if CPU use goes above a certain level, or the amount of free memory is below a certain level. Be sure to look at other options for remote server actions: full access

Figure 4
Viewing
Performance Data
Via Server Manager's
Performance Counters
Function



to Computer Management for the server, RDP, Windows PowerShell, configure NIC Teaming, reboot, and more. These are in addition to the standard tools on the Server Manager Tools menu, which lets you launch familiar tools such as Active Directory Users and Computers, Hyper-V Manager, and so on.

A big change to Server Manager in Server 2012 is the capability to remotely add and remove roles and features. To do this, select the *Add or Remove Roles and Features* action from the Manage menu. If you're performing a role-based or Remote Desktop Services installation, the next choice is the server on which you want to perform the add/remove (Figure 5).

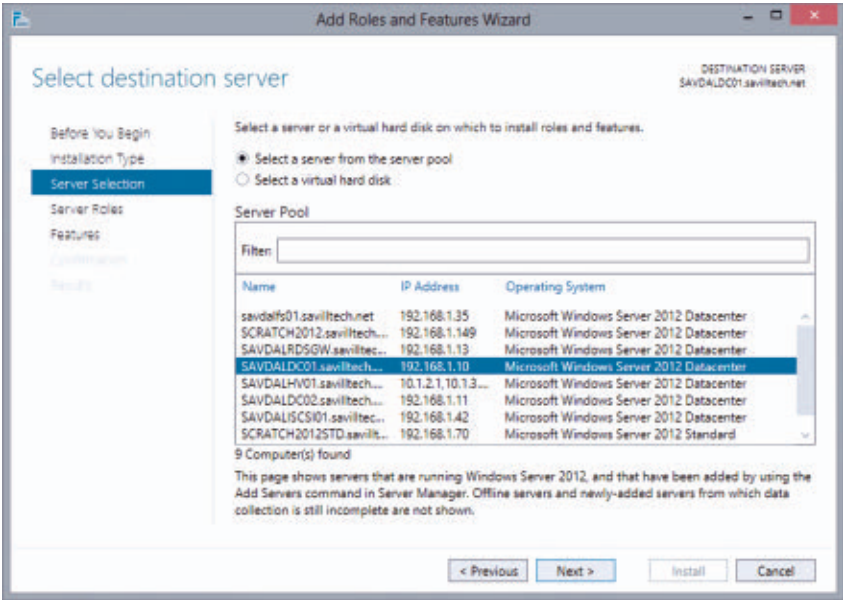


Figure 5
Selecting the Target Server in Server Manager for Add/Remove Operations

(I should note here that RDS has its own scenario-based deployment methodology accessible through Server Manager that enables you to deploy entire RDS environments across hundreds of servers—and based on best practices—from a single wizard with only a few clicks of a mouse.)

There are three things to note about this functionality:

1. You can select only one server at a time for the add/remove operation. If you want to add or remove roles/features on multiple servers, you must use PowerShell; for example:

```
Invoke-Command -ScriptBlock {Install-WindowsFeature -Name  
XPS-Viewer} -ComputerName server1,server2,server3
```

2. You cannot add/remove on Server 2008 and Server 2008 R2 boxes. Although it's possible to remotely manage these servers, this functionality doesn't extend to role/feature add/remove because the base OS doesn't support it.
3. Notice in Figure 5 that not only can you add/remove roles/features to a server, you can also select a virtual hard disk to manage offline—which is a great capability (you also can designate a virtual hard disk to manage offline from PowerShell as part of an automated provisioning process).

Manage roles and features. In addition to showing information (such as performance, events, BPA results, and installed components) from managed services, Server Manager also manages key aspects of the core OS, including certain roles and features. Storage is a great example. Server Manager not only exposes the volumes on a system and lets you configure features such as data deduplication, it also lets you manage storage spaces, iSCSI, NFS, and advanced share configurations. Some roles, such as Remote Desktop Services and IP Address Management, expose their configurations through Server Manager.

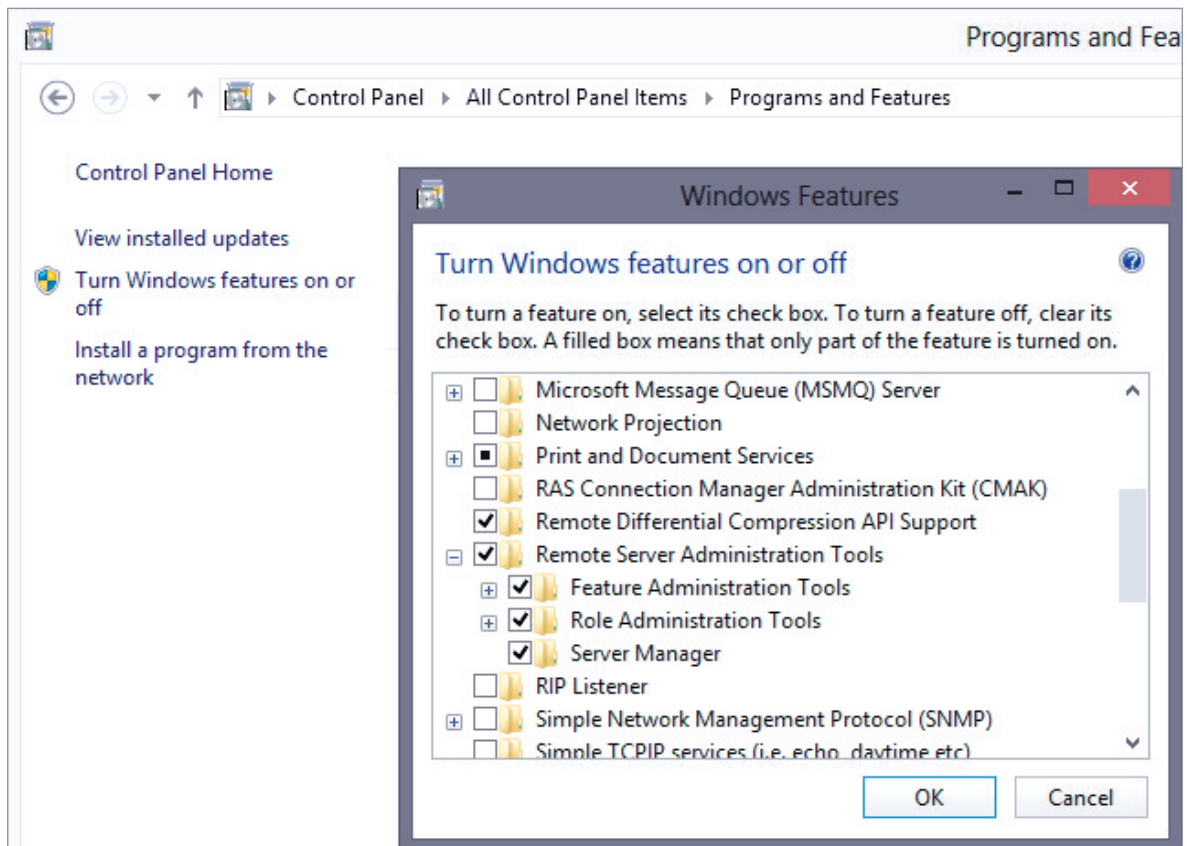
A common question regarding Server Manager is “How do you share a Server Manager configuration between users or different computers?” All servers and custom server groups monitored by a Server Manager instance are stored in an XML configuration file. It's therefore possible to move this configuration (which is stored at `%appdata%\Microsoft\Windows\ServerManager\ServerList.xml`) between users and computers.

This is all great; *however*, one of the first things I said in this article is Server 2012 installations should use the Server Core mode, which has no graphical interface and no management infrastructure available locally. So how is Server Manager being used? The answer is Server Manager runs on your Windows 8 desktop. Server Manager, along with the rest of the Server 2012 administration tools, is available as part of **RSAT, which is available in both 32- and 64-bit versions**.

Once the tools are installed on Windows 8, you enable them via the Control Panel *Programs and Features* applet. Run the *Turn Windows features on or off* action and enable Server Manager in the RSAT section (Figure 6). From this dialog box, you can enable other role and feature administration tools as well.

Figure 6

Enabling Server Manager on a Windows 8 Installation After Installing Remote Server Administration Tools (RSAT)



Video



John Savill
demonstrates Server
Manager in Windows
Server 2012



I step through this process and present a general overview in the accompanying video.

Once you have Server Manager installed on a Windows 8 desktop, you can realize the nirvana of Server 2012 management and deployment. All Server 2012 deployments run Server Core under the remote management of a Server Manager instance running on Windows 8.

RSAT for Server 2012 is available only on Windows 8. RSAT for Server 2012 doesn't support Windows 7 or even earlier client OSs; this is consistent with previous versions of Windows servers. RSAT for Server 2008 R2 ran on Windows 7 only, RSAT for Server 2008 ran on Windows Vista only, and Administration Tools Pack for Windows Server 2003 R2 ran on Windows XP only. You get the idea.

So what do you do if you don't have Windows 8 in your organization, and even you (as the IT department) can't adopt it? My recommendation is to set up a Server 2012 Remote Desktop Session Host and publish Server Manager. Then, administrators on Windows 7, Windows RT, and non-Microsoft platforms (basically, anything that supports published applications over RDP) can execute the published application, which appears to run locally on the machine because it

integrates seamlessly with the desktop. I created a [video](#), “Managing Windows Server 2012 from Windows 7,” that covers this procedure.

Ultimate Remote Management with PowerShell

Up to this point, I’ve discussed the new remote graphical management tooling. In the largest environments, however, where you’re talking about true automation, you don’t want a graphical interface. You want to use a command-line interface (CLI) and scripts. And for Windows, there really is nothing more powerful than PowerShell, which brings together a CLI and scripting environment. PowerShell is object-centric, allowing objects (rather than basic strings) returned by one command to be passed to another command. This provides for great capabilities with often minimal amounts of code because the original object is maintained.

Every Server 2012 capability is exposed through PowerShell by more than 2,400 cmdlets (the name for native PowerShell commands) just for the base OS; and the reality is that many things (such as Hyper-V network virtualization and advanced Storage Space configurations) can be achieved only through PowerShell and aren’t available in the OS graphical management tools. PowerShell integrates with Windows Remote Management to provide remote server management. PowerShell 3.0 (which is part of Server 2012 and installed as part of WMF 3.0 for Server 2008 and Server 2008 R2) also offers a powerful workflow capability that lets you initiate tasks on remote systems without having to stay connected.

Nearly all Microsoft products have PowerShell support, and so do most solutions that integrate with Windows, including those from hardware partners such as Cisco and NetApp. This means it’s possible to manage not only Windows with PowerShell, but also a complete IT infrastructure.

An introduction to PowerShell is beyond the scope of this article; however, I want to draw your attention to a few ways to get started quickly:

1. Use the Windows PowerShell Integrated Scripting Environment (ISE). It features IntelliSense, which helps complete cmdlet names and parameters as you type. ISE also has a command pane that lets you search for commands. The command pane provides a form interface that shows all available parameters and which ones are mandatory, and it creates commands for you.
2. Active Directory Administrative Center uses PowerShell behind the scenes. The PowerShell Command window displays the commands that accomplish the actions you perform in the graphical interface. This means you could graphically perform a series of actions and then incorporate the PowerShell commands into your scripts.
3. If you're experimenting, take advantage of the WhatIf parameter. Use a test box to limit damage as you're learning; however, the WhatIf command shows what would happen if the command were executed without making changes.

Remote Management Methodology

The new Server Manager in Server 2012 really helps organizations move to a true remote management methodology, letting you manage multiple machines simultaneously. It still might be a stretch for many IT groups to move away from just RDPing into each server and managing them “locally.” However, as more and more servers run Server Core, and once Server Manager’s multi-server management benefits become apparent, there’ll be a greater shift in management, especially as the 2012 Server Manager also provides for management of Server 2008 and Server 2008 R2. Of course, applications still must run on servers to support remote management, but this is also becoming more common and necessary as vendors add support for Server Core. And don’t forget about PowerShell, which really is an administrator’s best friend. ■

Use Active Directory Claims for Windows Server 2012 File Service Access Control

New claims and conditional expressions support in Windows Server 2012 file services simplifies file server access control

Windows Server 2012 Active Directory Domain Services (AD DS) added support for **claims** as part of the new Dynamic Access Control feature. Dynamic Access Control is a large and complex technology that lets you apply centralized governance, classification, and information protection on Server 2012 file servers. What has been largely overlooked, however, is that you can use claims right away on these file servers to give you new flexibility without waiting for a Dynamic Access Control implementation.

Security Groups

Using security groups to control access to files and folders is as old as the Windows NT Server file system. By adding groups to an ACL, then controlling group membership, this form of “crowd control” has enabled administrators to dynamically manage access to their resources without constantly manipulating the ACLs. Add the capability to nest these groups (e.g., to provide worldwide access to a folder, you might add NA_Sales, SA_Sales, EUR_Sales, and APAC_Sales global groups to the Sales domain local group), and you have a powerful and flexible access control model.



Sean Deuby

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel's core Directory Services team. He's been a Directory Services MVP since 2004.



Email



Twitter

It's not without its problems, however. First, security group lifecycle management is one of the most neglected areas of Active Directory (AD) hygiene. Everyone wants you to create groups in a hurry, but no one (except AD admins) seems to care about when it's time to eliminate them. As a result, most large AD domains have tens of thousands of security groups—many of them empty—with no plan to clean them up.

Second, you can run into trouble when you build a universal/global/domain local group access control architecture and don't account for the possibility of acquisitions or divestitures that may require multi-forest access.

Finally, group nesting, when combined with poor group lifecycle management, can lead to extremely complicated access control architectures. I think all admins would be surprised and dismayed to learn how many circular group references (a set of nested groups that refers to one of the groups in the set) they have in their environment.

Much of this complexity is reflected in a user's access token. The AD access token, stored in the Privilege Attribute Certificate (PAC) field of the user's Kerberos ticket-granting ticket (TGT), contains the user's SID and the SIDs of the security groups of which the user is a member. If, due to complex nested groups, the user is a member of many groups, that person will have large tokens that can lead to resource access problems.

One of the main reasons nested groups are used in this model is because they are the only architectural way to provide an AND operation. For example, consider a US-based program, classified secret, in a multinational corporation. One requirement is that only the HR people on the project have access to the folder that contains program employees' salary information. Further, because it's classified secret by the corporation, only full-time, US, HR employees can have access. Using only security groups, the access control configuration might look something like the diagram in Figure 1.

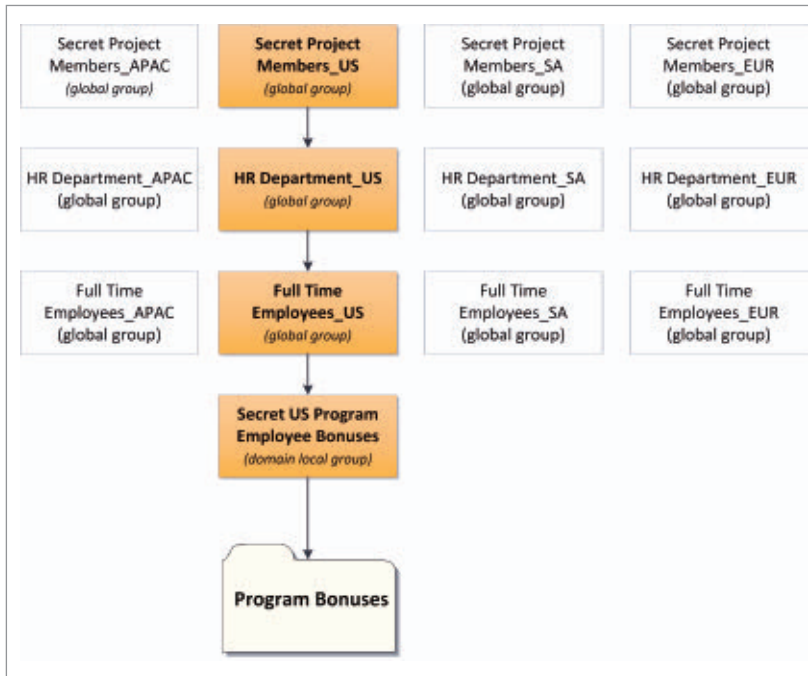


Figure 1
Nested Groups to
Enable “AND” Access
Conditions

This scenario calls for the creation of 12 groups to AND three conditions together to provide access. Multiply this situation hundreds of times in an international corporation, and you see the problem.

Claims and Conditional Expressions

The access control model in Figure 1 uses only security groups, but it doesn’t use an important repository of up-to-date corporation information: Active Directory itself. In most companies, AD receives much of its user attribute information directly from HR databases, so you know its data is up-to-date. But unless you’re using an external engine like a metadirectory service, these changes won’t be reflected in security group membership.

As I describe in my article [“Enable Claims Support in Windows Server 2012 Active Directory,”](#) Server 2012 AD DS supports exposing user and device attributes to File Services as claims. With this capability configured, you can add attribute values (such as department

name or number) to a file or folder's ACL to control access to it. While I don't ever see it supplanting security groups for access control—AD attribute classes are tightly defined and not easily created or modified—using AD attributes in the form of claims to supplement and potentially simplify group-based access control can be a great benefit to administrators. Next I explain how you would use the hybrid model for this example.

Two of the three conditions in this example, Full Time Employees_US and HR Department_US, can be replaced by claims. Look at Full Time Employees_US. I can break it down into two AD attributes. First, an employee's country can be represented by the two-digit attribute countryCode (00 would be a typical choice for the United States). And second, the attribute that best represents employee status would be employeeType. The latter attribute is string-based, so you can define it how you'd like. This is both good and bad; you'll want to keep it as simple as possible and (in my opinion) humanly readable. In other words, values of 0, 1, 2, and 3 are all good, but no one looking at the attribute can tell what the values represent. Instead, consider representing the values simply but with some clue as to their meaning: "F" or "FTE" for full-time employee, "T" for temp or "C" for contract, and "V" for vendor. Of course, this is English-centric, but I'm describing a US-based company.

You could describe the condition like this:

```
user.countryCode=00 AND user.employeeType=FTE
```

For HR Department_US, I can simplify things a little more. If I make the assumption that if the expression evaluates true (i.e., they are US and full-time employees), then if they are members of the HR department they will be members of the US HR department. To represent this with an AD claim, I can use the department attribute:

```
user.department=HR
```

The resulting representation of this hybrid access control architecture would look like the diagram in Figure 2.

You could even argue that the Secret US Program Employee Bonuses domain local group is redundant (though standardized access practices might dictate it should remain). Using conditional expressions, I reduced the number of groups necessary to represent this access condition by two thirds.

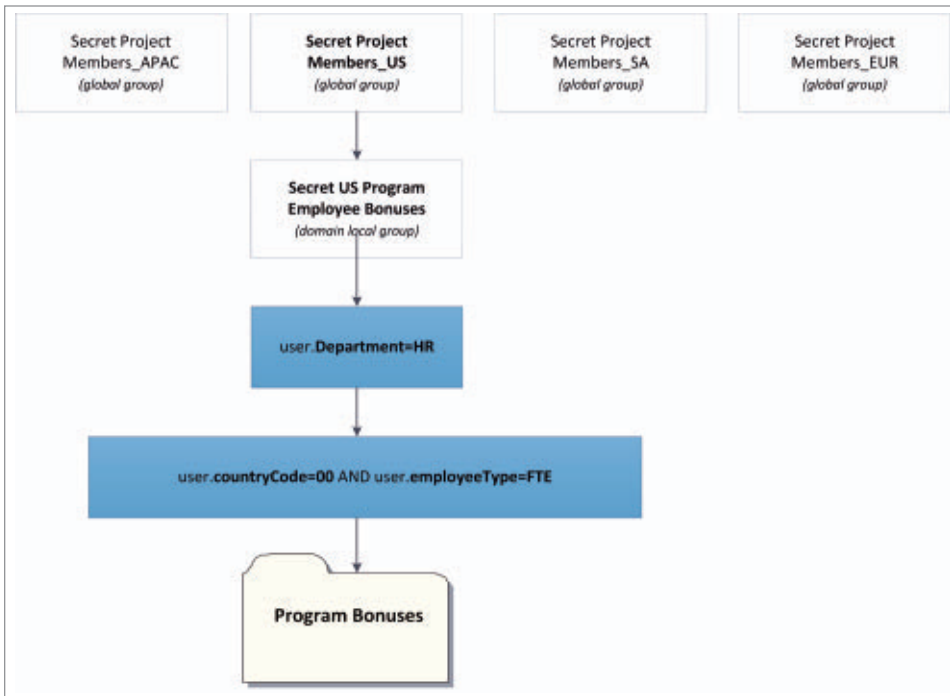


Figure 2

Hybrid-Groups Claims
Access Control

Example of How to Set Hybrid Permissions

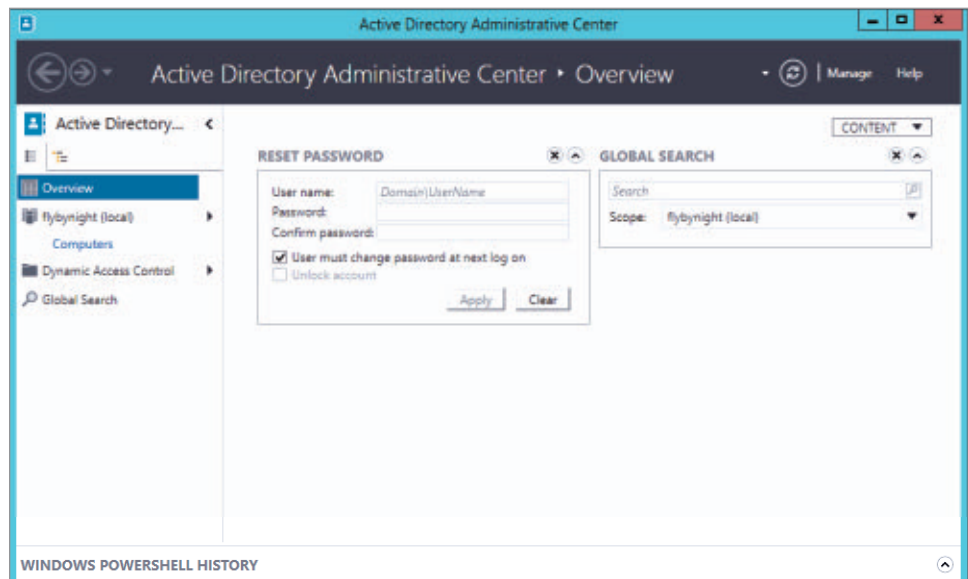
The following is a step-by-step example of how you would implement this hybrid set of permissions in a fictitious Fly By Night Software domain, flybynightsoftware.com. First, create the Secret US Program Employee Bonuses domain local group, and then the Secret Project Members_US global group. Add Secret Project Members_US into Secret US Program Employee Bonuses, and add this group into the Program Bonuses folder ACL. Next, build the claims-based controls.

Before you create claim types, you must enable claims for your domain via the process I outlined in “[Enable Claims Support in Windows Server 2012 Active Directory](#).” Then you must create claim types for the following user attributes:

- employeeType
- countryCode
- department

To do this, launch Active Directory Administrative Center (Figure 3).

Figure 3
Active Directory
Administrative Center



There are at least three ways to launch ADAC, but if you have a command prompt open, just type *DSAC*. In the left content pane, choose the arrow next to Dynamic Access Control and then select Claim Types (Figure 4).

The Tasks pane appears on the right; it gives you the ability to create new claim types. Claim types also appears under Dynamic Access Control in the left content pane.

Select New, Claim Type from the Tasks pane. This launches the Create Claim Type dialog box (Figure 5).

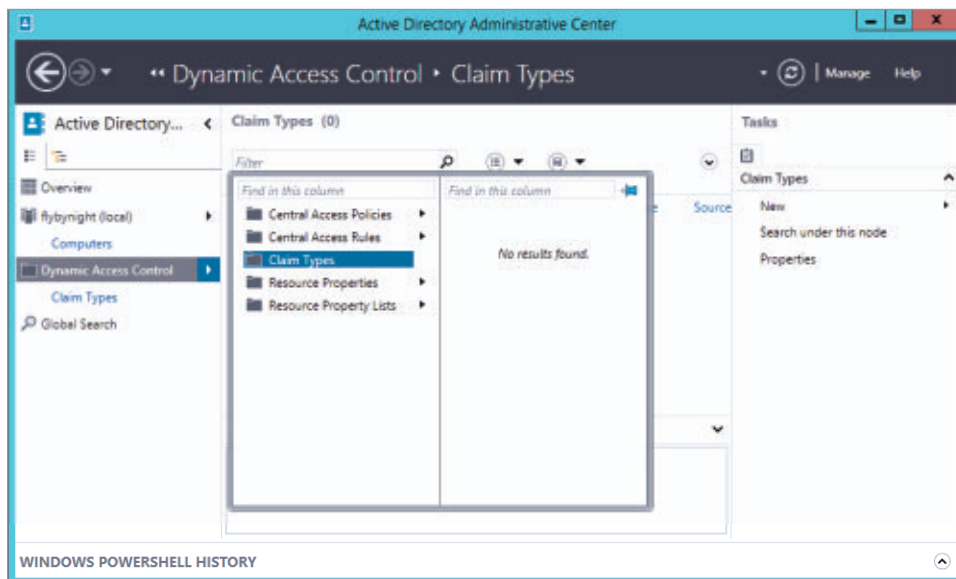


Figure 4
Selecting Claim Types

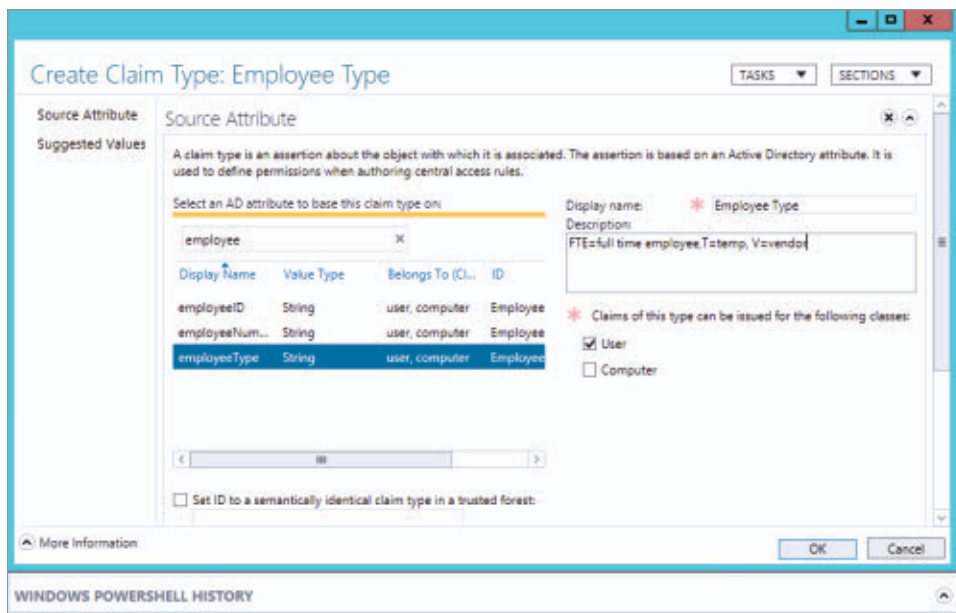


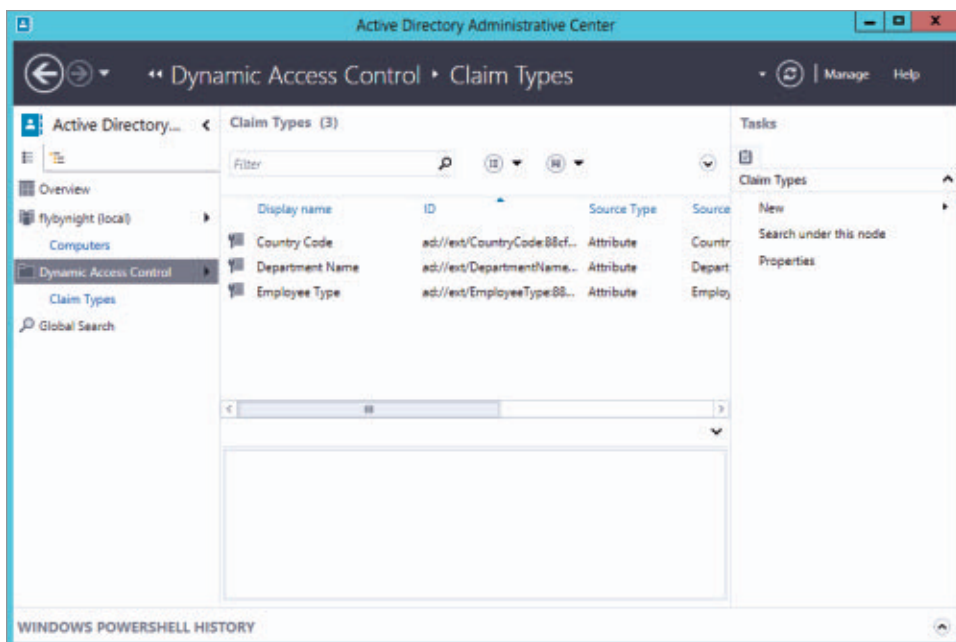
Figure 5
Creating a Claim Type

Enter *employee* in the Source Attribute filter field, as Figure 5 shows. This automatically narrows down the available attributes to three; choose *employeeType* from the list. Note that the display name

defaults to the display name of the attribute. If you want to make it a bit friendlier, change the display name to “Employee Type.” You also can update the description to include the different possible values of this attribute.

Don’t use line breaks in the Description field; using Return or hitting the Enter key will simply create the claim type before you’re ready. When you’re ready, click OK to create the claim type. Repeat this process for countryCode and department (see Figure 6 for an example).

Figure 6
All Claims Types
Created



Now you’re ready to configure the Program Bonuses folder with this mix of security groups and claims. Look at the folder’s properties, then the Security tab, then advanced permissions, and choose Add. Four layers deep into the File Explorer dialog boxes, you’ve arrived at the Permission Entry dialog box (Figure 7), which creates an access control entry (ACE) that you can add to the folder’s discretionary ACL (DACL).

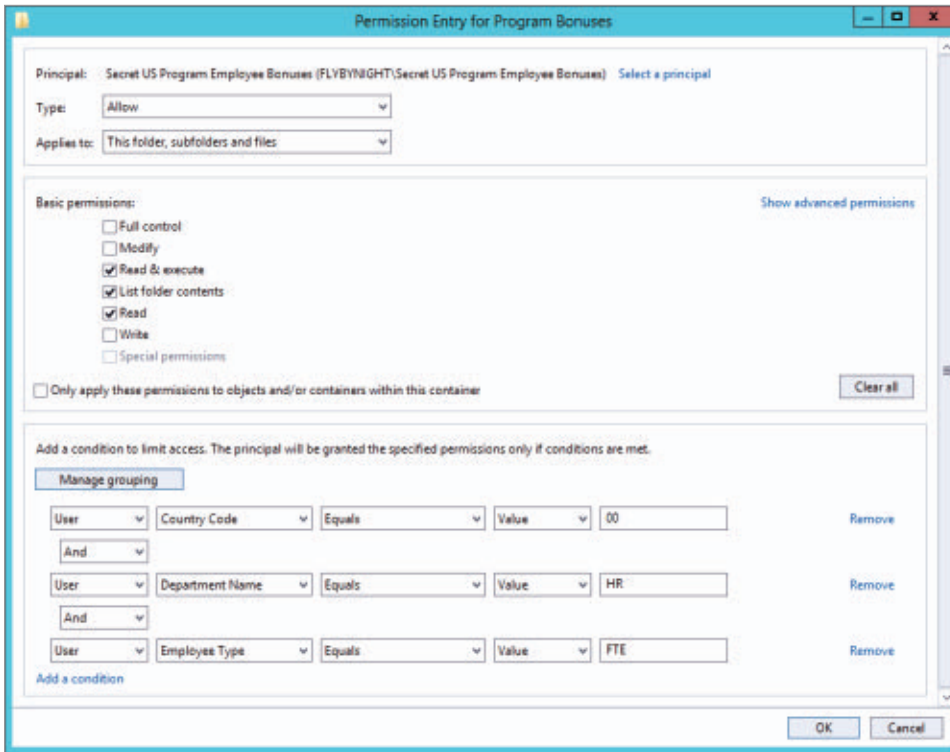
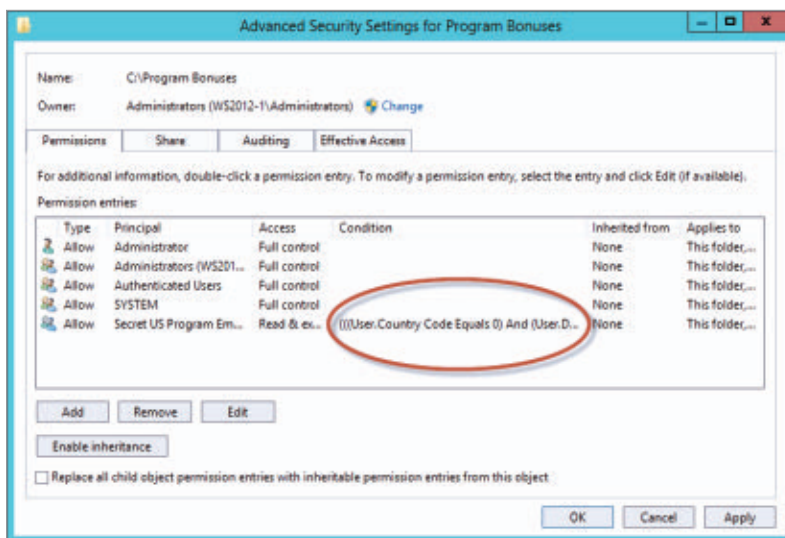


Figure 7
Adding Conditional
Expressions to the
Access Control Entry

There are three main sections in the Permission Entry dialog box. The first section lets you select a security principal (in this case, Secret US Program Employee Bonuses) to add to the ACE. Next is the type of access to be granted. The third section is new. This section lets you add conditions to limit access to the resources. (If you've worked with Apple's iTunes, this section looks familiar to how you build a smart playlist based on the attributes of your music files.) In Figure 7, I've added the three claims-based conditions that are necessary to meet the example's access requirements. You can see in Figure 7 how simple configuring conditional access is once you've made the appropriate claims available in the Permission Entry dialog box. Once you click OK, these conditions are displayed in the Conditions column of the Advanced Security Settings dialog box, which you can see in Figure 8.

Figure 8
ACL Showing
Conditional Expression
Access Control Entry



Claims-Based Access Control Caveats

Claims-based access control does have some management and governance concerns you should consider. In a group model, you're careful about who has rights to three of the four CRUD operations (create, read, update, delete) that change membership in a group. The same concerns exist for claims—you must pay attention to who has rights to populate, update, and remove AD attribute values. But because until now there has been little need for it, AD attribute governance in most companies isn't nearly as mature as group governance. If you want to use claims-based access control, you need to pay attention. Is this attribute populated by a “convenience” feed from another system versus a secured feed? Who in the HR department has the rights to populate the employeeType, countryCode, and department attributes? What if a user with this access should change the format of one of these attributes? Would it break some of your claims-based access control configurations in production? Would you receive advance warning about it?

You also need to be certain your AD data is clean and consistent within a domain and across domains if you have a multi-domain

forest. You can't count on a claim working consistently if the data it relies on isn't consistent. You also must be careful regarding the impact claims have on the Kerberos ticket-granting ticket (TGT). If you enable a claim type for a broadly populated attribute such as `employeeType`, that claim value will go into the security token of every user in the company. It isn't much in and of itself, but don't go crazy making populated AD attributes in claims without thinking through the ramifications. (If an attribute is only populated for some users, however, only those users will have it added to their security token.)

If you want to use claims-based access control for a special application or group, you'll run into an available attributes problem: Most attributes are already defined for a specific use. Unless you want to repurpose existing attributes (e.g., `telexNumber`: anyone remember a telex?) so that you can have a batch of unused, general purpose attributes, then you'll have to add your own attribute classes to AD by performing a schema extension.

Still, none of these caveats are insurmountable; you simply need to take your time and think through the situations in which claims-based access control will benefit your directory and business.

Experience in Claims

Implementing this sort of conditional access scenario is a good way to gain experience in the area before you begin something like a full-blown [Dynamic Access Control](#) test or pilot. It's worth repeating that claims-based access control isn't intended to replace the well-used group model. Rather, claims-based access control is a good addition to groups in situations where achieving the required access with groups alone would be too complicated. ■

FREE Newsletters!

**Not your average
Newsletters!**

WinInfo Daily UPDATE

Paul Thurrott covers the entire Windows universe with reviews, commentary, analysis, and tips. Delivered daily.

Windows IT Pro UPDATE

Windows industry news, products, FAQs, tips, and resources for IT professionals. Delivered weekly.

Cloud & Virtualization UPDATE

Get the latest news, blogs and analysis to help you determine your organization's cloud and virtualization strategy. Delivered weekly.

Exchange and Outlook UPDATE

News, strategies, products, and developments in Exchange Server and Outlook messaging. Delivered weekly.

Security UPDATE

Learn about Windows security risks, attacks, and how to fix or avoid them. Includes security alerts! Delivered bi-weekly.

Dev Pro UPDATE

Topics for Microsoft platform developers: ASP.NET, .NET Framework, Silverlight, mobile, and SQL Server development. Delivered weekly.

SQL Server Pro UPDATE

The latest news, products, and developments for SQL Server DBAs and developers. Delivered weekly.

SharePoint Pro UPDATE

SharePoint for IT professionals and developers – weekly tips, news, and how-to's. Delivered weekly.

**subscribe today at
windowsitpro.com/manage-newsletters**

WindowsITPro

PowerShell Basics: Console Configuration

How to customize the PowerShell console

The [Windows PowerShell](#) console provides an easy-to-use environment for creating and running PowerShell commands as well as generating script files that you can run at a later time. The more you work in that environment, the more likely you'll want to customize the console to meet your individual development style. For example, you might want to expand the buffer size or change the font and background colors. To configure these and other types of settings, PowerShell provides several methods that are easy to use. You can set the console properties directly from the command window, run PowerShell commands that configure these settings, or add scripts to your PowerShell profile so that the settings are automatically applied at startup.

Note that it's also possible to modify the registry to configure the PowerShell console, but a full discussion of this approach is beyond the scope here. Keep in mind that modifying the registry is neither the most flexible nor the easiest method for customizing the PowerShell console. Furthermore, it needs to be approached with great care.

Configuring the PowerShell Console Properties

When you first launch PowerShell, the console's appearance is determined by the default property settings or by the settings defined within a shortcut and passed in as arguments to the PowerShell executable. For example, if you run PowerShell by double-clicking the powershell.exe file directly, it opens with the default property settings. The same goes for launching PowerShell from a shortcut you created for the executable or launching PowerShell from the Windows command prompt. In most



Robert Sheldon

has written numerous books and articles about Windows technologies, database systems, business intelligence, scripting, enterprise operations, and consumerization. His books include *Beginning MySQL* (Wiley) and *SQL: A Beginner's Guide* (McGraw-Hill).




Email



Website


Figure 1
Launching the
PowerShell Console
from the
Executable File



cases, you'll see a small command shell window with a black background and gray font, as shown in Figure 1.

If you launch PowerShell from the Start menu, you get a somewhat different picture. In this case, the console window is larger with a blue background and gray font, as shown in Figure 2. It's different

Figure 2
Launching the
PowerShell Console
from the Start Menu



because the PowerShell installation process sets up the Start menu shortcuts with the configuration settings necessary to modify the console environment.

Regardless of how you launch PowerShell, you can modify the console's appearance through its property settings. To access those settings, click the PowerShell icon in the top-left corner of the console window and click Properties to open the Properties dialog box.

The Properties dialog box includes four tabs—Options, Font, Layout, and Colors—each of which contain configuration settings that you can modify as necessary. Figure 3 shows the Options tab. Here you can configure settings related to cursor size and command history. You also have two important editing options: QuickEdit Mode and Insert Mode. The QuickEdit Mode option lets you use your mouse to copy and paste commands in the PowerShell console. The Insert Mode option inserts new text into a line rather than overwriting it.

Figure 4 shows the Font tab, which includes options specific to the font used in the PowerShell console. Although these options are somewhat limited, you do have a few choices for size and style.

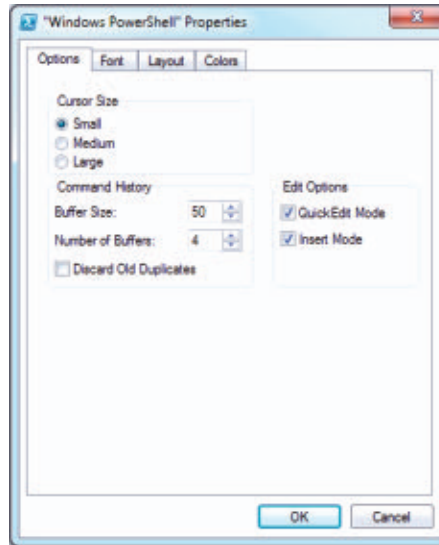


Figure 3
Configuring
PowerShell Options
in the Console's
Properties

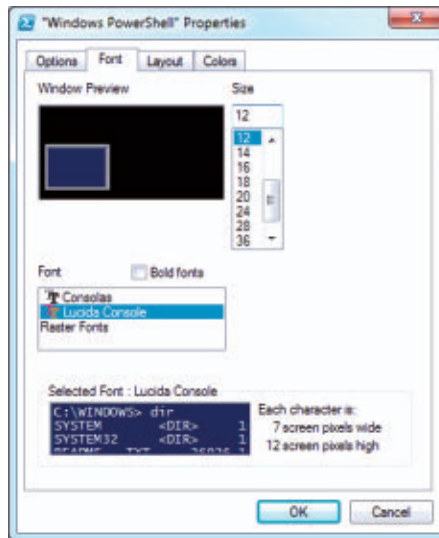



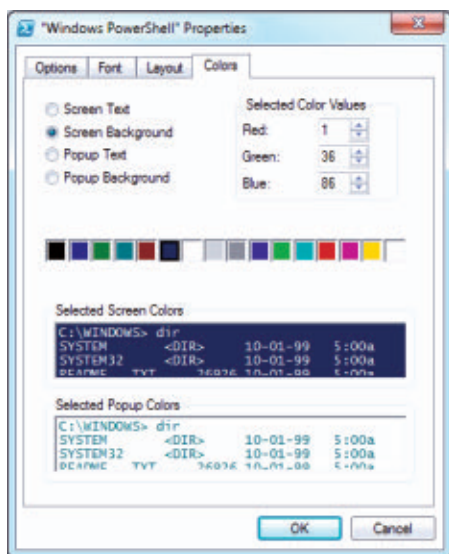
Figure 4
Configuring Font
Settings in the
Console's Properties

Figure 5
Configuring Layout
Settings in the
Console's Properties



and refer to them as needed by simply scrolling up or down. The window size determines the width and height of the actual console window. Again, this can't be larger than the buffer size. The window position determines where the console should be placed on your desktop, based on the top-left corner—or you can let the system position the window for you.

Figure 6
Configuring Color
Settings in the
Console's Properties



The Layout tab lets you choose the buffer size, window size, and window position, as shown in Figure 5. The buffer size determines the width (number of characters) and height (number of lines) in your buffer. Your buffer size can be the same size or larger than your window size, but it can't be smaller. Many administrators find it handy to increase the buffer size, particularly the height. That way, they can preserve many more commands during a session


Figure 6 shows the Colors tab, where you can choose your text and background colors. The easiest way to choose a color is to select a category option (e.g., Screen Text), then click a color in the color bar. If you know the specific Red Green Blue (RGB) value associated with a color, you can enter those numbers, but be aware that this will affect other items that share the same base color, resulting in some

rather bizarre effects. Luckily, you can view your changes in the preview windows to see what impact your changes are having on the console.

After you've configured the different properties, click OK to close the Properties dialog box. Your changes will be applied immediately. Figure 7 shows what the PowerShell console looks like if you enlarge it, change the background to white, and change the text to black.

Your changes will be preserved as long as you restart PowerShell from the same location from which you originally launched it. For example, if you launch PowerShell from the Start menu,

Figure 7
Examining a
Reconfigured
PowerShell Console



change the property settings, and relaunch PowerShell from the Start menu, your changes will be preserved. However, if you then launch PowerShell from a shortcut you created for the executable, your changes won't be reflected. To make your changes persist regardless of how you launch PowerShell, you have to take other steps, the first of which is to learn how to script your configuration settings.

Scripting the Console Configuration

PowerShell is first and foremost a scripting environment. Although PowerShell is often referred to as an interactive management console, you're still primarily writing and running scripts, even if they do interact with other systems. With PowerShell, you can script just about anything, including settings that affect the PowerShell console. Being able to control your settings through scripts provides an easy way to apply the same settings to multiple instances of the PowerShell console, whether they're on the same computer or different computers. You can save your settings in a script file and run it whenever you want the settings applied to your workspace. Or better still, you can save your settings in a profile file (which is covered in the next section) so that the settings are applied whenever you open PowerShell, no matter whether you open it by clicking the executable, a shortcut for the executable, or a shortcut on the Start menu. Scripting also lets you to access settings not available through the console's properties, such as the colors used for warning messages.

To script your PowerShell settings, it helps to have a basic understanding of how objects are handled in PowerShell. Objects form the basis for much of the scripting you do in PowerShell, even if it's not obvious that's what you're doing. Objects provide a structure for representing data within PowerShell. The structure is made up of properties, methods, and other members that you can access during your PowerShell session.

For example, when you run the `Get-Host` cmdlet, PowerShell returns an object that provides details about the PowerShell environment, such as the name, version, and information related to the shell's configuration. One of the members of the `Get-Host` object is the `UI` property, which is a special type of property associated with its own object, derived from a Microsoft .NET Framework class. The `UI` object, in turn, includes the `RawUI` property, which provides access to the specific console properties.

When working with objects in PowerShell, it's often best to assign them to a variable so you can easily access the object's members. As it turns out, PowerShell provides a built-in variable—`$host`—for accessing the `Get-Host` object. That means you can use the `$host` variable to access the `UI` and `RawUI` properties.

The `RawUI` property is a special type of property that's associated with its own object, just like the `UI` property. You access the console's properties through the `RawUI` object. Let's look at an example to help make sense of how all this works. The following command creates a variable named `$console` and assigns an instance of the `RawUI` object to that variable:

```
$console = $host.UI.RawUI
```

Notice that the `RawUI` object is accessed by first specifying the `$host` variable, then specifying the `UI` property, followed by specifying the `RawUI` property. By assigning the `$host.UI.RawUI` command to the `$console` variable, the variable is created as a `RawUI` object type, which gives you access to the console's properties so that you can configure them.

For example, the following commands set the `ForegroundColor` property (i.e., the text) and `BackgroundColor` property of the `RawUI` object:

```
$console.ForegroundColor = "black"  
$console.BackgroundColor = "white"
```

As you can see, you need to specify only the `$console` variable, followed by the property name. You then follow this with an equals sign (=) and the new color, enclosed in double quotes. In this case, the foreground color is being set to black and the background color is being set to white.

If you ran these two commands, you probably noticed that the changes were immediately applied to the console, leaving you with a rather odd-looking window because only the most recent lines reflect the changes to the background and foreground colors. The easiest way to make your screen presentable is to run the following command to clear the screen and start with a clean prompt:

Clear-Host

Now let's look at another RawUI property: `BufferSize`. As the name suggests, this property lets you set the buffer's width and height. However, this is another one of those special properties associated with its own object, so the best strategy is to define a variable to hold the object. You can then access the properties from that variable, as shown in this example:

```
$buffer = $console.BufferSize
$buffer.Width = 130
$buffer.Height = 2000
$console.BufferSize = $buffer
```

First, you create the `$buffer` variable to hold the `BufferSize` object. Then, you use that variable to set the `Width` and `Height` properties, similar to the way in which the foreground and background colors were set previously. However, you must take an additional step, which is to assign the settings in the `$buffer` variable to the actual `BufferSize` property of the RawUI object.

The same holds true if you want to set the size of the window itself, as shown in this script:

```
$size = $console.WindowSize
$size.Width = 130
$size.Height = 50
$console.WindowSize = $size
```

In this case, you first create a variable named `$size` to hold the object associated with the `WindowSize` property. You then set the `Width` and `Height` properties, just as you did for the `BufferSize` object. Finally, you assign the `$size` variable to the `WindowSize` property of the `RawUI` object. Note that when setting the window size, it can't be larger than the buffer (as mentioned earlier) and it can't be larger than what your system will support. If you receive an error that the window width or height can't be more than a specified size, adjust your code accordingly.

As handy as the `RawUI` object is, it doesn't contain all the properties that you can set. The `Get-Host` object also supports the `PrivateData` property, which itself is associated with an object. The `PrivateData` object includes a number of properties specific to the font and background colors used for console responses such as error and warning messages.

To simplify calling these properties, you can assign the `PrivateData` property to a variable, then call that variable, as seen in this script:

```
$colors = $host.PrivateData
$colors.VerboseForegroundColor = "white"
$colors.VerboseBackgroundColor = "blue"
$colors.WarningForegroundColor = "yellow"
$colors.WarningBackgroundColor = "darkgreen"
$colors.ErrorForegroundColor = "white"
$colors.ErrorBackgroundColor = "red"
```

First, you need to create a variable named `$colors` to hold the `PrivateData` object. Then, you can use the variable to access several of the object's properties in order to set their colors. This is basically the same approach used to set the `ForegroundColor` and `BackgroundColor` properties of the `RawUI` object. You define simple assignments, without having to assign the `$colors` variable back to the `PrivateData` property, as you did with `BufferSize` and `WindowSize`.

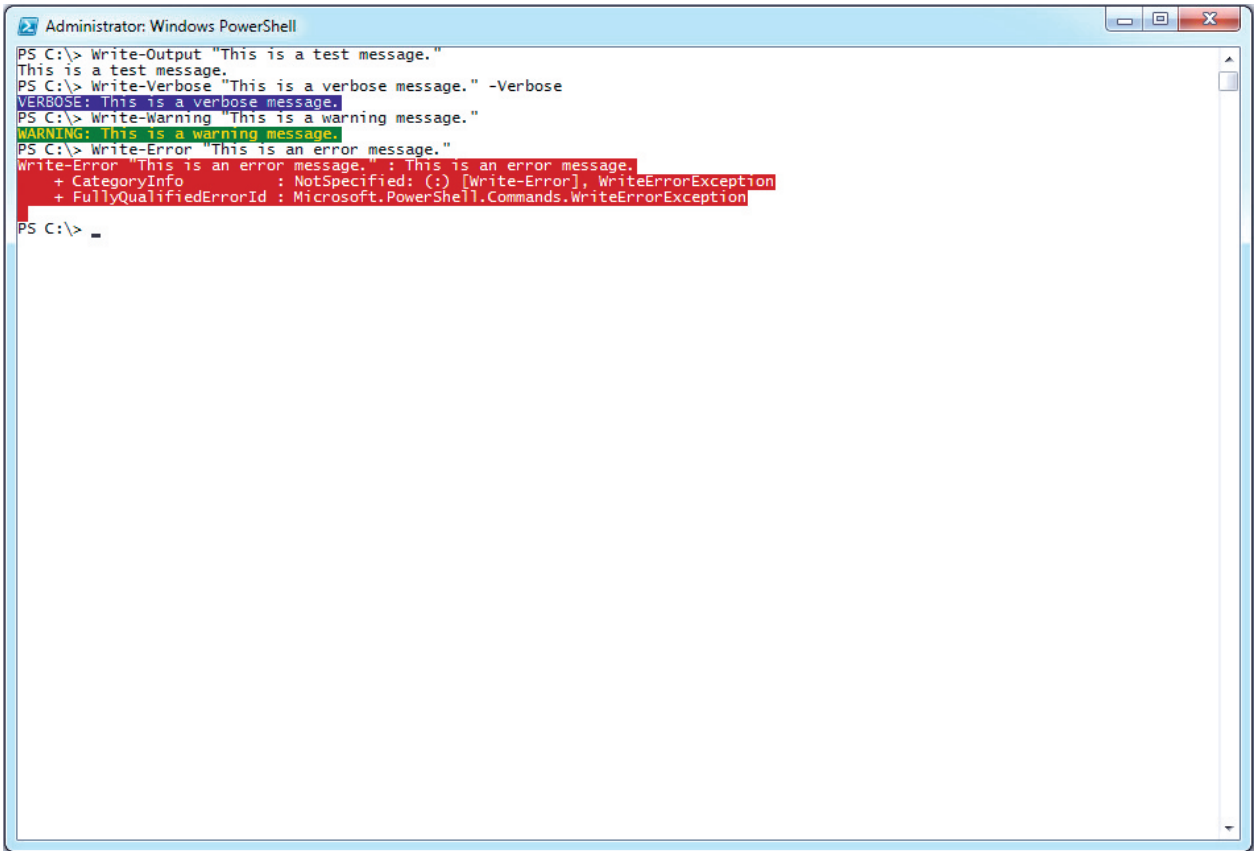
After you've defined all your properties, run the `Clear-Host` cmdlet again:

`Clear-Host`

Your screen should now look similar to what you see in Figure 7. Now let's test the new format that you created by running a few example commands. The following example commands generate several types of messages that should be displayed in the colors you configured previously:

```
Write-Output "This is a test message."
Write-Verbose "This is a verbose message." -Verbose
Write-Warning "This is a warning message."
Write-Error "This is an error message."
```

For the most part, the cmdlets used in these commands are self-explanatory. The `Write-Output` cmdlet generates a regular message in the default background color and font. The other three cmdlets generate the types of messages reflected in their names. Figure 8 shows what your console should look like after you generate these test messages. Notice that the first message uses a white background with black text, but the other messages reflect the colors that you assigned to the properties associated with the `PrivateData` object.



```
Administrator: Windows PowerShell
PS C:\> Write-Output "This is a test message."
This is a test message.
PS C:\> Write-Verbose "This is a verbose message." -Verbose
VERBOSE: This is a verbose message.
PS C:\> Write-Warning "This is a warning message."
WARNING: This is a warning message.
PS C:\> Write-Error "This is an error message."
Write-Error "This is an error message." : This is an error message.
+ CategoryInfo          : NotSpecified: (:) [Write-Error], WriteErrorException
+ FullyQualifiedErrorId : Microsoft.PowerShell.Commands.WriteErrorException
PS C:\> _
```

Adding the Console Configuration Script to Your Profile

Being able to script your console settings makes it easy to apply those settings whenever you need them. You can save them to a script file and call that file when necessary. Or you can cut and paste the script into the console and run the commands that way. In either case, you must access your script file whenever you want to apply the configuration settings to the console. A better approach is to save the script to one of the PowerShell profile files.

PowerShell supports four types of profile files and executes them in a specific order upon startup. However, a full discussion of PowerShell profiles is beyond the scope here. So, let's just set up the profile file that applies to the current user.

Figure 8

Testing the Fonts
in the PowerShell
Console

To find where the file should be located, you can use the built-in variable `$profile` to return the path and filename used for the file. To do so, simply run the command:

```
$profile
```

The `$profile` variable will return a fully qualified filename, such as `C:\Users\Administrator\Documents\WindowsPowerShell\Microsoft.PowerShell_profile.ps1`. However, just because the `$profile` variable points to this file, it doesn't mean the file exists. Therefore, the next step is to run the command:

```
Test-Path $profile
```

The `Test-Path` cmdlet checks for the existence of the file. If it exists, the command returns `True`; otherwise, it returns `False`. If the command returns `False`, you should run the following command to create the file before you do anything else:

```
New-Item -path $profile -type file -force
```

The `New-Item` cmdlet creates the file based on the argument passed in with the `-path` parameter, which in this case is the `$profile` variable. The `-type` parameter indicates that the item being created is a file (as specified by the "file" keyword) rather than another type of item. The `-force` parameter tells the cmdlet to create the file, no matter how Windows might balk.

After you've finished creating your file, you can rerun the `Test-Path` command to ensure that it returns `True`. Once you know that the file is in place, you can easily edit it by running the following command:

```
notepad $profile
```


This command opens the blank profile file in Notepad. You can then add whatever configuration script you want to include. If you add the configuration code that I previously discussed, the file will look similar to the one shown in Figure 9. (Note that you can download this configuration code by clicking the Download the Code button.) However, notice that the script in the figure also includes the command:

```
Set-Location C:\
```

This command merely tells PowerShell to set the startup folder to C:\ whenever the settings are applied. You can specify whatever



Download the code

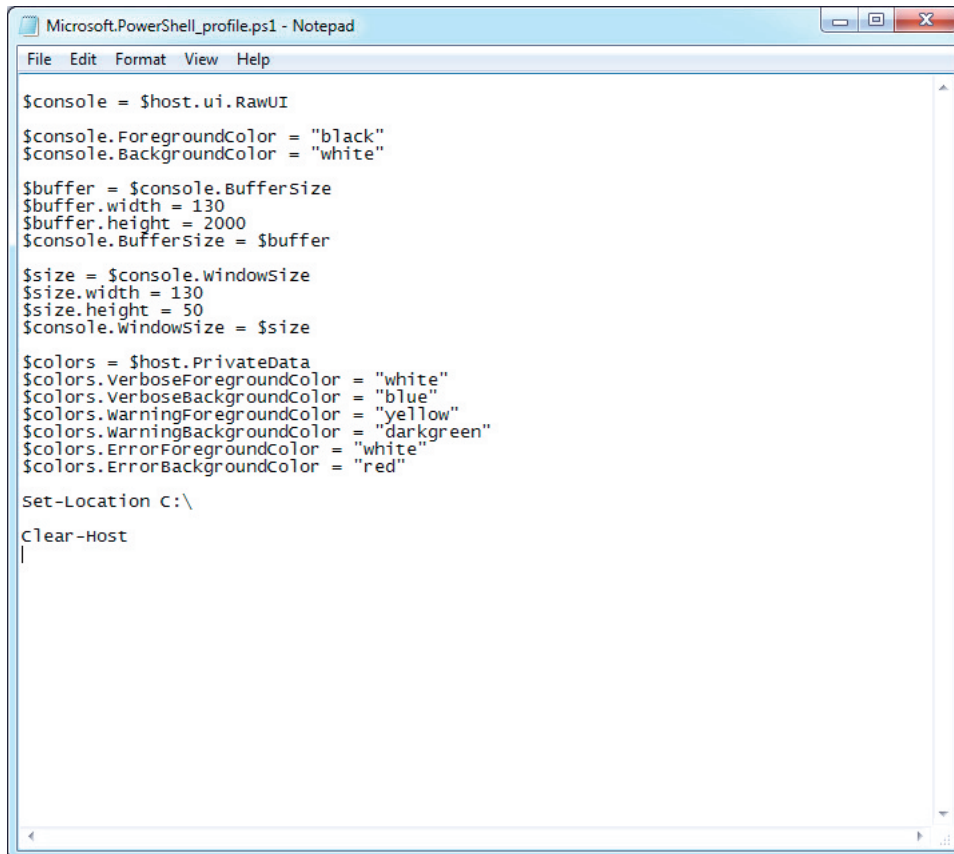


Figure 9

Editing the Profile File
in Notepad

folder you want. You can also specify the startup folder in a shortcut or within a PowerShell session. However, adding this command to your profile file helps keep all the configuration settings in one location. Note that you can also include additional commands in your profile file, such as defining variables that you might want to use in each session.

After you're satisfied that the profile file contains everything it needs, you simply save it and close Notepad. You must then relaunch PowerShell for the settings to be applied. However, you can relaunch PowerShell from any location. The same profile file will be applied in each case. In fact, it's even possible to save your profile file to a different location, such as a jump drive or network share, so that you can apply your settings to other instances of PowerShell, regardless of which computer you're working on. This approach also lets you share a common profile file among a team of users. Be sure to check the PowerShell documentation for more information about working with profile files.

If you haven't set up PowerShell to run script files (which includes profile files), you might receive an error the first time you relaunch PowerShell. Figure 10 shows what the error looks like. Basically, it's telling you that your execution policies have not been properly configured to support script files.

To permit scripts to run on your local system, run the command:

```
Set-ExecutionPolicy RemoteSigned
```

This command lets you run local scripts or remote scripts that have been digitally signed by a trusted publisher. Keep in mind, however, that the execution policy controls what scripts can and can't be run on your system. An incorrect setting could jeopardize your system's security. Be sure to carefully read the information about execution policies in PowerShell's documentation before making any changes to your system.




Figure 10
Receiving an Execution
Error in PowerShell

Console Customization

PowerShell provides a fair amount of flexibility when setting up your console. You can set the console's properties, run scripts that configure those settings, or modify your profile file to apply the settings whenever you launch PowerShell. Clearly, adding the necessary configuration script to your profile file will provide you with the most flexibility and least maintenance over the long haul. The profile file will persist your script, so you don't have to reapply your settings each time you start PowerShell. This approach also provides more options than the console's properties, and you don't have to be concerned about whether you start PowerShell from the Start menu or another location. Given the powerful scripting environment that PowerShell provides, there's no reason you shouldn't make the console itself as useful and comfortable as possible. ■

Using Microsoft System Center 2012 Configuration Manager for Updates

How to set up a Patch Tuesday update solution



Kent Agerlund

is a Microsoft System Center 2012 Configuration Manager MVP who works as senior System Center architect, trainer, event speaker, and author. For the past four years, he has been on the road with his Mastering System Center 2012 Configuration Manager class.

Email



One of the many features supported by [Microsoft System Center 2012 Configuration Manager \(SCCM 2012\)](#) is software updates. For any business, being and staying compliant is of the utmost importance. When setting up a software update solution, it's really important that you start with first things first—and the first thing is planning.

Planning for Software Updates

An important part of the planning process for SCCM 2012 is developing criteria that you can use to determine when you have reached an acceptable compliance level for updates. Without that information, it'll be difficult for you to know when you have to spend additional time tracking noncompliant devices. Table 1 shows sample compliance criteria for workstations. Table 2 shows sample compliance criteria for servers.

Table 1: Example of Compliance Criteria for Workstations

Update Severity Level	Success Criterion for Week 1	Success Criterion for Week 3	Success Criterion for Week 5
Extremely critical (zero day exploit)	90%	95%	99.5%
Critical	50%	80%	95%
Security	50%	75%	90%

Table 2: Example of Compliance Criteria for Servers

Update Severity Level	Success Criterion for Week 1	Success Criterion for Week 3	Success Criterion for Week 5
Extremely critical (zero day exploit)	99%	99%	100%
Critical	50%	80%	100%
Security	50%	75%	100%

Your planning must also include what to do when the criteria aren't met. Another important part of the planning process is determining what updates you want to apply and how often. Most organizations create unique deployments for each Patch Tuesday. (Microsoft releases security patches the second Tuesday of each month, which is often referred to as Patch Tuesday.) So, I'll walk you through how to set up a Patch Tuesday deployment. First, though, you need to be familiar with the components in a software update solution.

Understanding the Software Update Components

The software update feature in SCCM consists of eight components. Most of them only need to be created once, and the creation of the other components can be automated. After the components are created, approving and deploying monthly updates can take less than 10 minutes. The components and the recommended strategy for how often they should be created are as follows.

Software update point. A software update point is a Windows Server Update Services (WSUS) server controlled by SCCM. Unlike a standalone WSUS solution, clients don't download or install updates directly from a software update point. The only data downloaded by the client from a software update point is the update metadata. In SCCM 2012, only one software update point is supported, but multiple software update points are supported in SCCM 2012 SP1. You only need to install this component once.

Deployment package. A deployment package is like any other package in SCCM, except that it contains only the software update binary files. The client downloads only the required updates. As a result, deployment packages can contain a mix of updates from multiple OSs. In SCCM 2012 SP1, a client can fall back to Windows Update if the requested update isn't available in a deployment package. You should create a new deployment package twice a year.

Software update groups. A software update group is a group of updates that can be deployed to devices. They can also be used to track update compliance. A software update group can be created automatically using the Automatic Update Rule feature or manually by selecting the updates. You should create a new software update group every month for a Patch Tuesday deployment.

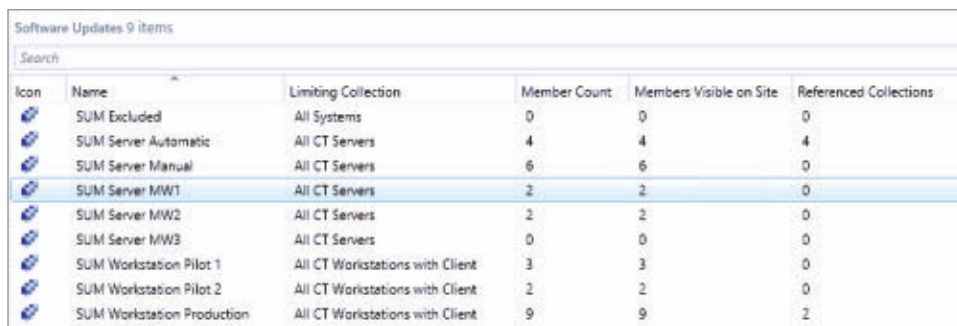
Deployments. The deployment is a child object of a software update group. Like any other deployment, it contains information about the installation purpose, schedule, and user experience (e.g., whether to restart the computer after an update if needed). The Automatic Deployment Rule will create the first deployment. All other deployments in the software update group will need to be created manually. You'll have to create a number of deployments each month.

Software update templates. Software update deployments can be controlled by the use of templates. You should create one template for each unique deployment scenario. Here are some sample templates you might consider creating:

- Pilot 1: All computers that participate in the first test deployment.
- Pilot 2: All computers that participate in the second test deployment.
- Workstation Production: All workstations that aren't excluded from patch management.
- Server Automatic: All servers in which the installation and restart will be performed automatically but controlled through maintenance windows.
- Server Manual: All servers in which the installation and restart will be performed manually.

Each template needs to be created only once.

Collections. A collection is a group of targets for a deployment. Each collection is created only once. You'll have at least one collection per template. Figure 1 shows some sample collections. Collections containing the letters MW all have a configured maintenance window. The Referenced Collections column specifies the number of referenced collections. A referenced collection is a collection that's either included or excluded in another collection. The SUM Excluded collection contains devices that won't be part of the update process.



Icon	Name	Limiting Collection	Member Count	Members Visible on Site	Referenced Collections
	SUM Excluded	All Systems	0	0	0
	SUM Server Automatic	All CT Servers	4	4	4
	SUM Server Manual	All CT Servers	6	6	0
	SUM Server MW1	All CT Servers	2	2	0
	SUM Server MW2	All CT Servers	2	2	0
	SUM Server MW3	All CT Servers	0	0	0
	SUM Workstation Pilot 1	All CT Workstations with Client	3	3	0
	SUM Workstation Pilot 2	All CT Workstations with Client	2	2	0
	SUM Workstation Production	All CT Workstations with Client	9	9	2

Figure 1
Exploring Some
Sample Collections

Maintenance windows. A maintenance window is a collection attribute that defines when software can be installed and when computers are restarted. A device will apply maintenance windows from all the collections of which it is a member. You create a maintenance window once.

Automatic Deployment Rule. The Automatic Deployment Rule is a very powerful feature that lets you fully automate the software update deployment process. The rule contains information about the run time, what updates to download, where to store the updates, and whether the deployment will be automatically enabled. It's common to have a rule for Patch Tuesday and a rule for System Center Endpoint Protection updates. For each application, you need to create an Automatic Deployment Rule once.

Setting Up a Software Update Deployment

Now that you know about the software update components, I'll guide you through the steps needed to set up a software update deployment for Patch Tuesday. I'll show you how to create a collection (including a maintenance window), create an Automatic Deployment Rule, work with software update groups, and deploy the updates to production machines. I don't describe how to create the software update point. For information about its creation, see the [Configuring Software Updates in Configuration Manager](#) web page.

Creating a Collection

You always deploy software updates to a collection, so creating collections is an important part in setting up a software update solution. You can add members to a new collection three ways:

- You can use a *direct rule* to add members to a new collection.
- You can use an *include collection rule* to include members of another collection in the new collection.
- You can use a *query rule* to dynamically add members to the new collection. With this method, you need to specify a WMI Query Language (WQL) query.

You can use Windows PowerShell, a new feature in SCCM 2012 SP1, to create a collection and directly add members to it. For example, to create the SUM WRK Pilot I collection with the Active Directory (AD) group SUM_WRK_Pilot1 as a member, you'd follow these steps:

1. Click the Home tab in the SCCM 2012 administrator console and select *Connect via Windows PowerShell*.
2. In the PowerShell console, type

```
New-CMDeviceCollection -Name "SUM WRK Pilot1"
-LimitingCollectionName "All Systems"
```

and press Enter.

3. While still in PowerShell, run the command to add the AD group SUM_WRK_Pilot1 as a member, such as:

```
Add-CMDeviceCollectionQueryMembershipRule
-CollectionName "SUM Workstation Pilot 1"
-RuleName "SUM Pilot 1"
-QueryExpression "select SMS_R_SYSTEM.ResourceID,
SMS_R_SYSTEM.ResourceType,
SMS_R_SYSTEM.Name,
SMS_R_SYSTEM.SMSUniqueIdentifier,
SMS_R_SYSTEM.ResourceDomainORWorkgroup,
SMS_R_SYSTEM.Client from SMS_R_System where
SMS_R_System.SystemGroupName =
'Domain\\SUM_WRK_Pilot1'"
```

(Although this command wraps here, you'd enter it all on one line in the PowerShell console.)

You can also create collections using either the Create Device Collection Wizard or the Create User Collection Wizard. For more information about creating collections using these wizards, see the [How to Create Collections in Configuration Manager](#) web page.

After you create a collection and add members to it, you have the option to create a maintenance window for it. Although the SUM WRK Pilot I collection you just created doesn't need a maintenance window, here are the steps you'd follow if you wanted to create one for another collection:

1. Open the properties of the collection.
2. Select the Maintenance Windows tab.
3. Click the Yellow starburst icon and fill in the details specifying the schedule for the maintenance window.
4. Click OK to save the changes, and close the collection properties.

Creating the Automatic Deployment Rule

With the collection created, you can use the Create Automatic Deployment Rule Wizard to create the Automatic Deployment Rule for your Patch Tuesday updates. Here are the steps:

1. In the SCCM 2012 administrator console, navigate to the Software Library workspace.
2. Select Software Updates, and choose Automatic Deployment Rules. Click the Create Automatic Deployment Rule option on the ribbon to launch the Create Automatic Deployment Rule Wizard.
3. On the General page, which Figure 2 shows, specify Patch Tuesday in the Name field and a description in the Description field. In the Collection field, enter or browse to the SUM WRK Pilot I collection you created. For the *Each time the rule runs and finds new updates* option, select *Create a new Software Update Group*. Although adding updates to an existing software update group

Figure 2
Specifying the
General Information
for the Automatic
Deployment Rule

Create Automatic Deployment Rule Wizard

General

Specify the settings for this automatic deployment rule

Name: Patch Tuesday

Description: Download all Critical and Security updates released every Patch Tuesday.

Select a previously saved deployment template that defines configuration settings for this deployment. You can save the current configuration as a new deployment template on the Summary page of this wizard.

Template: Manage Templates...

Specify the target collection for the software update deployment.

Collection: SUM Workstation Pilot 1 Browse...

Each time the rule runs and finds new updates.

☐ Add to an existing Software Update Group

☒ Create a new Software Update Group

Choose whether to enable the deployment after this rule runs for the associated software update group. When this setting is not selected, you must manually deploy the software update group.

☐ Enable the deployment after this rule is run

< Previous Next > Summary Cancel

is useful when creating an Automatic Deployment Rule for End-point Protection definition updates, it's not useful for regular software updates. Here you'll create a new group every month. Otherwise, you'll end up having too many updates in the group. (A software update group has a limit of 1,000 updates.) Clear the *Enable the deployment after this rule is run* check box. Click Next.

4. On the Deployment Settings page, click Next.
5. On the Software Updates page, select the following filters and add the specified search criteria: Date Released or Revised: Last 1 month; Update Classification: "Critical Updates" OR "Security Updates"; Title: -Itanium. Note that the Title filter will prevent updates containing the word *Itanium* from being downloaded. Confirm that your page looks like the one in Figure 3, then click Next.

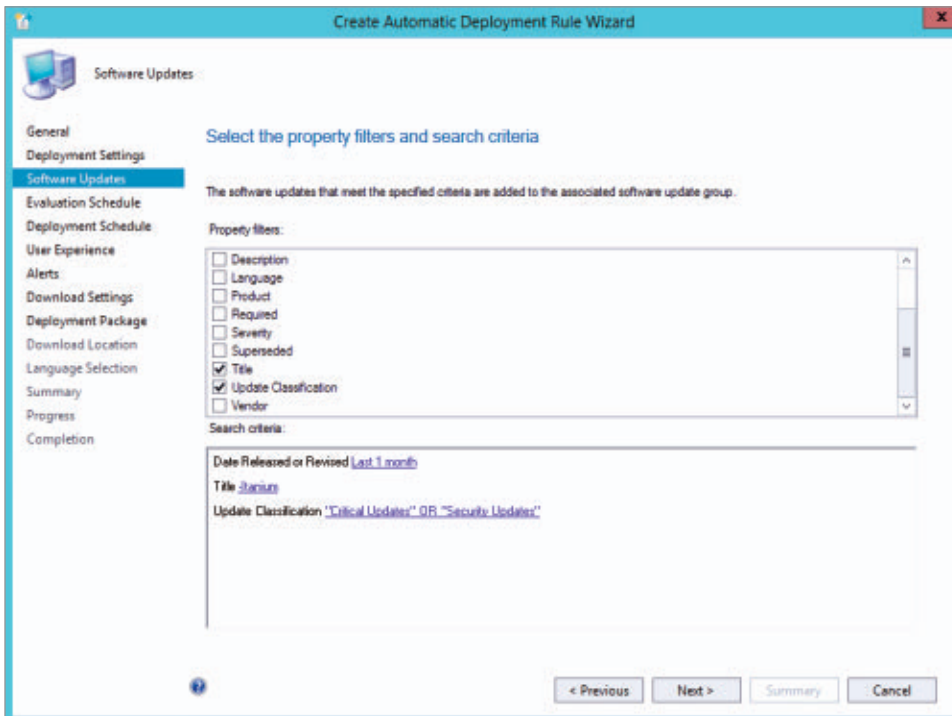


Figure 3
Specifying the Filters
and Search Criteria

6. On the Evaluation Schedule page, select *Enable rule to run on a schedule* and click the Customize button. Configure the rule to run the second Tuesday of every month at a time of your choosing. Click OK, then click Next.
7. On the Deployment Schedule page, configure the following settings. In the *Time based on* drop-down list, select *Client local time*. In the *Software available time* and *Installation deadline* sections, select *As soon as possible*. You don't have to worry about this deadline being too aggressive because this setting is applied only to the devices in your pilot group. For the production workstations, I recommend making the updates available two days prior to the company-decided deadline. Updates will start downloading in the background when they become available and will install when the deadline is reached. Click Next.
8. On the User Experience page, select *Display in Software Center and show all notifications* in the *User notifications* drop-down list. In addition, suppress the system restart on both servers and workstations, as shown in Figure 4. Click Next.
9. On the Alerts page, you can configure SCCM to send an alert when the compliance level drops below a certain percentage. To do this, select the *Generate an alert when the following conditions are met* check box. Then, in the *Client compliance is below the following percent* drop-down list, select 95. Finally, set the *Offset from the deadline* option to 35 days. This means that SCCM will generate an alert if the compliance level isn't at 95 percent 35 days after the specified deadline. Click Next.
10. On the Download Settings page, configure the following settings. Select *Download software updates from distribution point and install* as the deployment option for the preferred distribution point. Select *Download and install software updates from the fallback content source location* as the deployment option to use when updates aren't available on any preferred distribution point. Select the *Allow clients to share content with other clients*

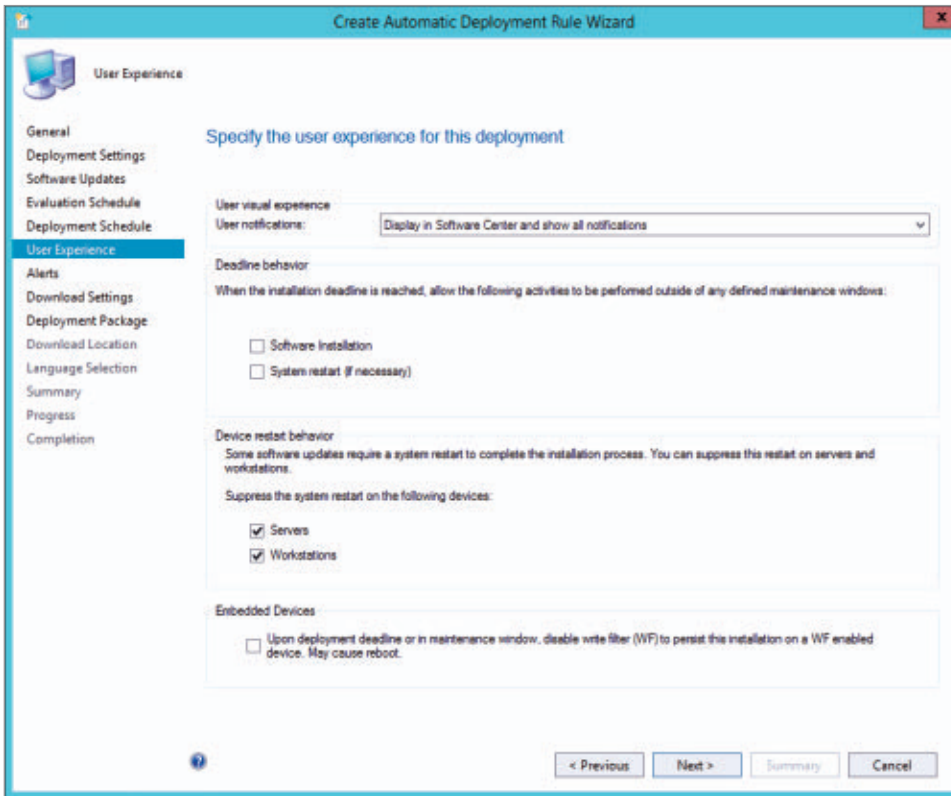
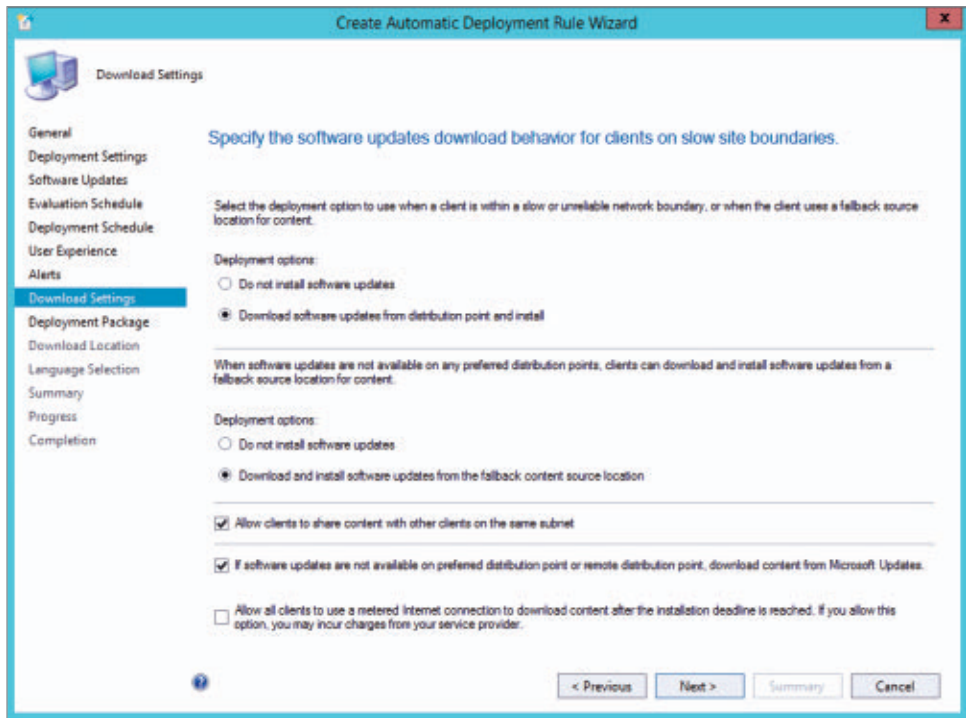


Figure 4
Configuring the User
Experience Settings
for the Automatic
Deployment Rule

on the same subnet check box. Select the *If software updates are not available on preferred distribution point or remote distribution point, download content from Microsoft Updates* check box. This is a new SP1 feature that allows clients to fall back and use Windows Update to download the content. The client will only download content for the updates you approved. After making sure your settings look like those in Figure 5, click Next.

11. On the Deployment Package page, you can either select an existing deployment package or create a new one. For this example, create a new one, specifying a name and description for it. In the Package Source field, enter or browse to the folder containing the software update binary files. Leave the sending priority at the default of medium. Click Next.

Figure 5
Specifying How to
Download the Updates



12. On the Distribution Points page, specify the distribution points or distribution point groups to which you want to distribute the package and click Next.
13. On the Download Location page, select *Download software updates from the Internet* and click Next.
14. On the Language Selection page, select the languages supported in your organization and click Next.
15. On the Summary page, click Save As Template. In the Save As Template dialog box that appears, type Pilot Deployment I in the Name field and click Save.
16. Click Next to have the wizard create the Automatic Deployment Rule. When it completes, click Close.

You'll now see the Patch Tuesday rule in the list of Automatic Deployment Rules. Manually run that rule by selecting it and clicking the

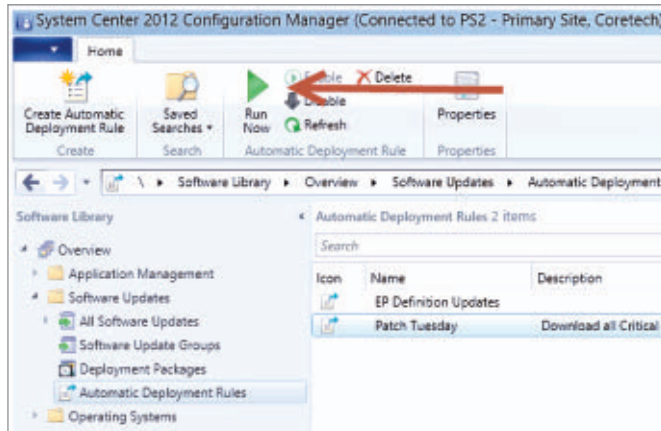


Figure 6
Running the
Automatic
Deployment Rule
Manually

Run Now option on the ribbon, as shown in Figure 6. Click Yes to start the process.

Working with Software Update Groups

The Patch Tuesday rule will now automatically create a new software update group every Patch Tuesday. What you need to do every month is rename the update group, remove any unwanted updates, and enable the pilot deployment.

To rename the update group and remove any unwanted updates, follow these steps:

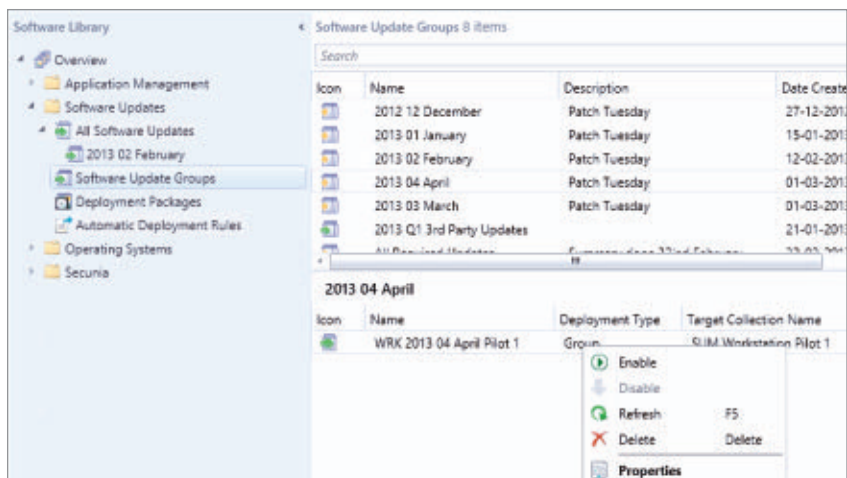
1. In the SCCM 2012 administrator console, go to the Software Library workspace. Navigate to Software Update Groups and verify that you have a new Patch Tuesday update group.
2. Rename that update group by right-clicking it, selecting Properties, and entering the new name. Naming standards are as important in SCCM as in any other management system. You'll be using the names when running reports and tracking update compliance. I recommend that you use a naming convention such as Year + Number of Month + Name of Month (e.g., 2013 04 April).
3. Right-click the update group and select Show Members. When you navigate down the different updates in the group, notice that the compliance statistics are updated.

4. Remove any unwanted updates from the update group by right-clicking the update, selecting Edit Membership, and choosing *Remove the update(s) from the shown update groups*.

At this point, it's time to enable the deployment of the Patch Tuesday updates to the workstation pilot group. Follow these steps:

1. Go back to the Software Update Groups workspace, select the renamed update group, and click the Deployment tab at the bottom of the window. Notice that you have a disabled deployment.
2. Right-click the deployment, select Properties, and change the name to something more descriptive by including the details about the collection and whether it's a pilot deployment (e.g., WRK 2013 04 April Pilot I).
3. Right-click the deployment and click Enable, as Figure 7 shows.

Figure 7
Enabling the
Workstation Pilot
Deployment



Deploying the Updates to Production Machines

After a successful deployment to your pilot group, you're ready to create the deployment for the production workstations. To do this, you use the Deploy Software Updates Wizard. Follow these steps:

1. Make sure the software update group is selected and click Deploy on the ribbon to launch the Deploy Software Updates Wizard.

2. In the Deployment Name field on the General tab, type the name of the deployment.
3. In the Collection field, enter or browse to your collection containing your production workstations. Click Next.
4. On the Deployment settings page, click Next.
5. On the Scheduling page, specify an installation deadline to determine when updates will be installed automatically. Click Next.
6. On the User Experience page, configure the following settings. In the *User notifications* drop-down list, select *Display in Software Center, and only show notifications for computer restarts*. In the *Device restart behavior* section, select the Workstations check box. Confirm that your page looks like the one in Figure 8, then click Next.

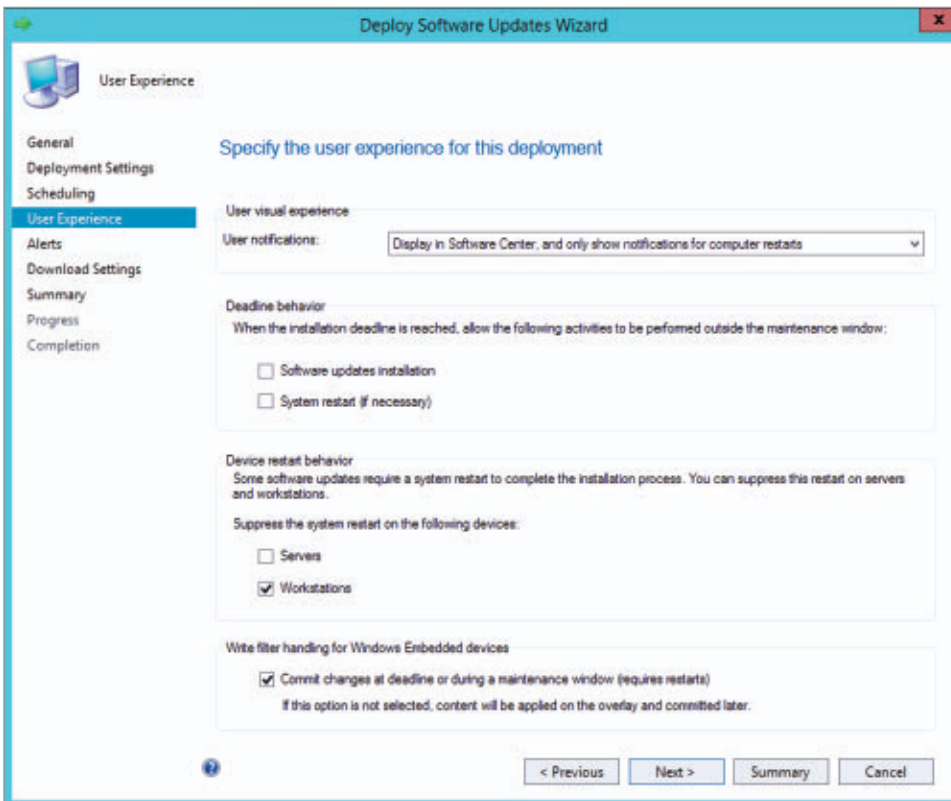


Figure 8

Configuring the User Experience Settings for the Workstation Production Deployment

7. On the Alerts page, select *Generate an alert when the following conditions are met* and click Next.
8. On the Download Settings page, select the option *If software updates are not available on preferred distribution point or remote distribution point, download content Microsoft Updates*. Click Next.
9. On the Summary page, click Save As Template, type Workstation Production, and click Save. Saving the settings as a template prevents you from having to go through these same steps every month.
10. Click Next.
11. Click Summary and Next to deploy the updates.

You might want to create additional Patch Tuesday deployments, such as a deployment for the servers that can restart automatically and a deployment for the servers that require a manual restart. Assuming that you already added some servers to a pilot collection and tested the Patch Tuesday updates against it, you can use the Deploy Software Updates Wizard to deploy those updates to the production servers.

For example, the following steps show how to deploy the Patch Tuesday updates to servers that can restart automatically using a pre-defined template:

1. Make sure the software update group is selected and click Deploy on the ribbon to launch the Deploy Software Updates Wizard.
2. In Deployment Name field on the General tab, type the name of the deployment, as shown in Figure 9.
3. Click the Select Deployment Template button and select a pre-defined template.
4. If needed, click any of the page links (e.g., Scheduling, Alerts) if you want to change elements of the deployment.
5. Click Summary and Next to deploy the updates.

Deploy Software Updates Wizard

General

Specify general information for this deployment

Deployment Name:

Description:

The following software update or software update group is included in this deployment.

Software Update/Software Update Group:

Select a previously saved deployment template that defines configuration settings for this deployment. Before you complete this wizard, you have the option to save the current configurations as a new deployment template.

Template has been applied: Server Automatic

Deploy this software update deployment to the following collection.

Collection:

< Previous Next > Summary Cancel

Figure 9
Specifying the General Information for a New Deployment Based on a Template

Keep It Simple

For many administrators, handling software updates is a complex process, but it doesn't have to be. By keeping it simple, as shown here, you'll likely get the job done quicker and gain a better understanding of the software update process. You should always start by defining how many deployments you need and defining a service level agreement (SLA) that's approved by management and is achievable by you. Once you have the deployments and SLA defined, SCCM 2012 is a great tool to ensure high compliance with a minimum of effort. ■

FAQ

Answers to Your Questions



John Savill




Jan De Clercq

Q: What URL do I type to configure SkyDrive Pro on my machine?

A: SkyDrive Pro works in a similar manner to the consumer SkyDrive but synchronizes work-based content with your organization's SharePoint installation (on-premises SharePoint 2013 or Office 365). When you first launch the SkyDrive Pro application, it will ask for the URL to use to configure synchronization. This URL should have been automatically configured using Group Policy or been given to you by your administrator. If you don't know it, use the process below when leveraging Office 365:

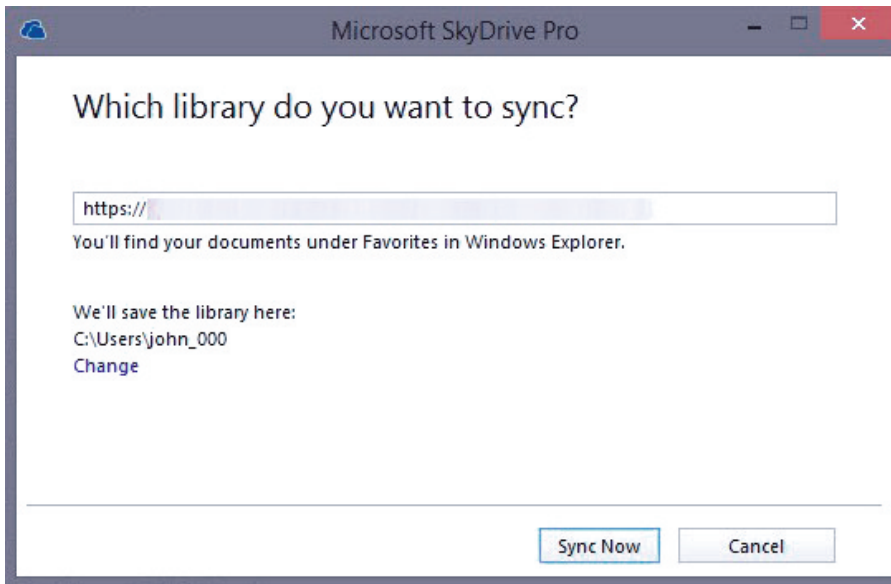
1. Log on to your Outlook Web Access site (<http://www.outlook.com/> < your company DNS >).
2. Select the SkyDrive area, and click the SYNC button in the top right corner (see Figure 1).

Figure 1
SYNC Button for
Microsoft SkyDrive Pro



3. A screen will be displayed, prompting you to set up easy access to your library (see Figure 2). Click the Sync Now button.
4. A dialog will be displayed asking you to confirm if a program can be opened on your computer. Click Allow.

Figure 2
SkyDrive Pro Screen
Prompting Library
Access Setup



5. The SkyDrive Pro application will launch, and the URL will be automatically completed along with the location to cache content to. Click the Sync Now button.
6. The synchronization will run in the background; a dialog box will inform you of the synchronization and give you the option to open the folder on your machine. You will then have completed configuring SkyDrive Pro.

—John Savill

Q: I'm trying to import an exported virtual machine (VM) from Windows Server 2008 R2 into Windows Server 2012 R2, but it's not working. What might the problem be?

A: If you try to import a VM into Windows Server 2012 R2 that you had exported by using Windows Server 2008 R2, the import process will say that Hyper-V didn't find VMs to import from the location you specified.

This is because Server 2012 R2 supports importing only from [Windows Server 2012](#) and Server 2012 R2. To import those VMs exported using Server 2008 R2, you will first need to import to a Server 2012 machine, then import those files into Server 2012 R2 (either by exporting from Server 2012 or just directly importing the files from Server 2012). Export isn't actually required to import with Server 2012 and later.

—John Savill

Q: Is it possible to join a Windows client computer to an Active Directory (AD) domain when the client computer is located in a branch office that has only read-only domain controllers (DCs) or when no network connection is available to the central site holding read-write DCs?

A: Yes, both are possible starting with [Windows 7](#) and [Windows Server 2008 R2](#), thanks to the djoin.exe command-line utility. You can use djoin.exe to provision a computer account in AD and export the account's AD security information to a text file. This text file can be moved and imported to the client computer, where an administrator then has to run djoin.exe from an elevated command prompt to effectively join the domain. The content of the text file can also be added to an unattended setup answer file to join a computer to the domain as part of the OS installation.

For example, suppose you want to join a Windows client computer named mywindowsclient to the AD domain mydomain.com. First, you need to create the AD computer account for mywindowsclient in mydomain.com and save the domain join data to a text file named offlinejoin.txt. This can be accomplished by running the following Djoin command from an elevated command prompt on a Windows 7 (or later) or Server 2008 R2 (or later) machine that can communicate with a read-write DC:

```
Djoin /provision /domain mydomain.com
/machine mywindowsclient /savefile c:\offlinejoin.txt
```

(Although this command wraps here, you'd put it on all one line in the command-shell window.) Next, you need to join mywindowsclient to the mydomain.com domain by running the following Djoin command from an elevated command prompt on the mywindowsclient machine:

```
Djoin /request0DJ /loadfile c:\offlinejoin.txt
/windowspath %systemroot% /localos
```

(Although this command wraps here, you'd put it on all one line in the command shell window.) You must then reboot mywindowsclient. When it comes back up, it'll be joined to the domain.

Note that you can also provision a computer's AD account against DCs running Windows Server 2008 or earlier by using the /downlevel switch in the first Djoin command. For more information about offline domain joins, see the TechNet article "[Offline Domain Join \(Djoin.exe\) Step-by-Step Guide](#)."

—Jan De Clercq

Q: How can I rename a large number of files by using Windows PowerShell?

A: I recently wanted to rename a portion of a large number of files. The best way I found to accomplish this was by using [PowerShell](#), which made it very easy to rename only a specific part of the file name. I used this command:

```
get-childitem *.wmv | foreach {Rename-Item $_ $_.Name.Replace
("oldstring","newstring")}
```

—John Savill

Q: To protect Active Directory (AD) user and computer objects from accidental deletion, how do I set the *Protect object from accidental deletion* property?

A: When the *Protect object from accidental deletion* property is enabled for an AD object, the object's permissions are automatically set to deny the deletion of the object by the built-in Everyone group. By default, it's enabled only in AD organizational units (OUs). When the property is set, it doesn't propagate down to child objects in the OU—it applies to the OU object only.

You can manually enable the *Protect object from accidental deletion* property on an AD user or computer object from the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in or the Active Directory Administrative Center GUI. Another way to enable the property is by running a Windows PowerShell command. With PowerShell, you can even simultaneously enable the property for multiple AD objects. For example, the following command enables it for all AD user and computer objects:

```
Get-ADObject -filter {(ObjectClass -eq "user")} |  
Set-ADObject -ProtectedFromAccidentalDeletion:$true
```

This command uses the `Get-ADObject` cmdlet to get any object with the value "user" in the `ObjectClass` attribute, which includes all AD user and computer objects. The command then uses the `Set-ADObject` cmdlet to enable the *Protect object from accidental deletion* property in those objects. `Get-ADObject` and `Set-ADObject` are part of the Active Directory Module for Windows PowerShell, installed by default on Windows Server 2012 or Windows Server 2008 R2 domain controllers (DCs). You can also install the module on non-domain DCs running Server 2012 or Server 2008 R2, as well as on computers running [Windows 8](#) or Windows 7. ■

—Jan De Clercq

Product News for IT Pros

Spiceworks 7 Provides Free Mobile Device Management

Spiceworks released Spiceworks 7, the latest version of its free IT management platform. Spiceworks 7 includes, for the first time, mobile device management (MDM) capabilities, as well as significant updates to its inventory and Help desk applications. Collectively, the new features can give you greater visibility and control over the on-premises, [cloud](#), and mobile technologies you use across your networks. Spiceworks 7 also includes updates to its inventory and scanning capabilities to help IT pros get a full picture of their environment, including the computers, servers, routers, switches, and applications running on their network. The inventory application lets IT pros customize their scanning and inventory experiences, and new scanning workflows help IT departments inventory hard-to-decipher devices in an intuitive way that makes the scanning and inventory process easier. These additions allow IT pros to easily control scanning across very large, complex network environments. For more information and to download Spiceworks 7, visit the [Spiceworks website](#).



Advanced Systems Concepts Brings Windows Azure Support to ActiveBatch IT Automation

Advanced Systems Concepts has become the first job scheduling/workload automation vendor to integrate with Windows Azure, Microsoft's platform for building, deploying, and managing applications and services in the cloud. The new capability enables IT departments to dynamically allocate Azure resources to match workload processing requirements in order to reduce IT operational costs. ActiveBatch 9 changed the workload automation landscape by allowing IT organizations to combine dynamic workload automation and



management with the power and flexibility of cloud computing. By automating the management of Azure instances within ActiveBatch, the product gives users an even more effective way to instantly allocate cloud resources to workload processing, when and where they're needed. The new ActiveBatch Extension for Azure allows users to provision and manage Azure instances within workflows that direct other applications, platforms, and process types, all from within ActiveBatch's centralized interface. To maximize efficiency, users can take advantage of ActiveBatch's Smart Queue capabilities to provision Azure instances automatically based on workload demands. For more information, check out the [Advanced Systems Concepts website](#).



IronKey Workspace USB Drives Now Certified for Windows To Go

Imation announced that the IronKey Workspace W500 has been certified by Microsoft for Windows To Go. The IronKey Workspace W500 features hardware encryption and a rugged metal enclosure, with enterprise-grade deployment and device-management options, including mass provisioning and IT management of the portable workspaces of mobile and remote users. The result is a fast, durable, and secure IronKey “Workspace PC on a Stick” USB drive that lets IT meet the needs of a flexible, mobile workforce with Windows To Go, and protect the organization with always-on hardware encryption. Features and benefits of IronKey Workspace W500 include military-grade security with hardware-based AES 256-bit encryption and strong authentication to help keep data safe and secure; “set and forget” ease of use for a provisioned, managed, and secured workspace infrastructure; a rugged metal casing that protects against physical damage, and sealed components to defend against tampering; superior speed, with more than five times the minimum read/write performance required for Windows To Go certified devices; and a range of capacity options for users' needs (32GB, 64GB, and 128GB). For more information, visit the [IronKey website](#).

Riverbed Delivers Enterprise-Ready Cloud Storage Appliances

Riverbed Technology announced that it has expanded its Whitewater cloud storage appliance family with the addition of new hardware models and upgrades to its OS. The new Riverbed Whitewater appliances and OS provide more capacity, faster ingest speeds, and more replication options. These features make the new Whitewater appliances a critical component for enterprises wishing to leverage the economical price and reliability of cloud storage options such as Amazon Glacier. Enhancements include new Whitewater model appliances with up to triple the cache of previous models and support of up to 14.4 petabytes of logical data. The Whitewater Operating System (WWOS) 3.0 also offers new features, including pairwise replication, which enables enterprises to replicate to an additional Whitewater appliance at a secondary location. In addition, enterprises can now leverage the 10Gb networking interface that dramatically improves ingest performance. WWOS 3.0 is a free upgrade for Whitewater customers. For more information, visit the [Riverbed Technology website](#).



Onvelop Launches on Windows 8 App Store

Onvelop—an enterprise mobility collaboration and communication platform that utilizes licensed Microsoft protocols to provide secure access to enterprise communication and collaboration tools such as SharePoint, Lync, and Office 365 from smartphones and tablets across multiple OSs such as Android, iOS, and Windows Phone—is now available in the Windows 8 App Store. Onvelop's user experience brings enterprise mobility and BYOD to business consumers without requiring the installation of any new software at the enterprise back end. It seamlessly integrates with existing infrastructures, allowing users to stay connected and fully able to collaborate with colleagues—all from a mobile device, in real time. Users simply enter their Office 365 logon credentials, and they're ready to work. For more information, visit the [Onvelop website](#).





Secunia Launches Secunia CSI 7.0

Secunia released the new version of its flagship solution, Secunia Corporate Software Inspector (CSI) 7.0, which introduces new vulnerability and patch management features to organizations worldwide. To help IT teams counter the threat of cybercrime, Secunia merges its in-house vulnerability expertise with a sophisticated patch management solution into CSI 7.0. The foundation of the product is a unique combination of vulnerability intelligence and vulnerability scanning, with patch creation and patch deployment integration. The solution integrates with Microsoft Windows Server Update Services (WSUS), [System Center 2012](#), and third-party configuration management tools for easy deployment of third-party updates, making patching a simple and straightforward process for all IT departments. For more information, check out the [Secunia website](#).



Acronis vmProtect 9 Targets Smaller VMware Environments

Acronis launched Acronis vmProtect 9, an efficient and easy-to-use VMware vSphere backup solution for small-to-midsized businesses (SMBs) and remote branch offices of large enterprises. With added support for [Microsoft SQL Server](#), [SharePoint](#), and Active Directory (AD), vmProtect remains the only solution capable of recovering data at different levels: hypervisor, virtual machine (VM), and application. Acronis vmProtect combines dozens of advanced technologies packaged in an easy-to-use product that installs in less than three minutes. The solution's accessible web UI is optimized for smaller vSphere environments. Unlike complex products with many features designed for large deployment sizes, Acronis vmProtect is focused on maintaining cost-effectiveness and ease of use. For more information, visit the [Acronis website](#). ■

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowsitpro.com

Support
Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
forums.windowsitpro.com

News
Check out the current news and information about Microsoft Windows technologies.
www.winsupersite.com

EMAIL NEWSLETTERS
Get free news, commentary, and tips delivered automatically to your desktop.

- Cloud & Virtualization UPDATE
- Dev Pro UPDATE
- Exchange & Outlook UPDATE
- Security UPDATE
- SharePoint Pro UPDATE
- SQL Server Pro UPDATE
- Windows IT Pro UPDATE
- WinInfo Daily UPDATE

RELATED PRODUCTS
Windows IT Pro VIP
Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.
windowsitpro.com/vip-premium-membership

SQL Server Pro
Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

Dev Pro
Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

SharePoint Pro
Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.
www.sharepointpromag.com

Advertiser Directory

Software Developer's Journal..... 1

Windows IT Pro..... 2, 23, 56

Vendor Directory

Acronis.....	96	Onvelop.....	95
Advanced Systems Concepts.....	93, 94	Riverbed Technology.....	95
Concur Technologies.....	23	Secunia.....	96
IronKey.....	94	Spiceworks.....	93

